arXiv:1402.1009v1 [math.OC] 5 Feb 2014

DYNAMIC PROGRAMMING SUBJECT TO TOTAL VARIATION DISTANCE AMBIGUITY

IOANNIS TZORTZIS*, CHARALAMBOS D. CHARALAMBOUS†, AND THEMISTOKLIS CHARALAMBOUS‡

Abstract. The aim of this paper is to address optimality of stochastic control strategies via dynamic programming subject to total variation distance ambiguity on the conditional distribution of the controlled process. We formulate the stochastic control problem using minimax theory, in which the control minimizes the pay-off while the conditional distribution, from the total variation distance set, maximizes it.

First, we investigate the maximization of a linear functional on the space of probability measures on abstract spaces, among those probability measures which are within a total variation distance from a nominal probability measure, and then we give the maximizing probability measure in closed form. Second, we utilize the solution of the maximization to solve minimax stochastic control with deterministic control strategies, under a Markovian and a non-Markovian assumption, on the conditional distributions of the controlled process. The results of this part include: 1) Minimax optimization subject to total variation distance ambiguity constraint; 2) new dynamic programming recursions, which involve the oscillator seminorm of the value function, in addition to the standard terms; 3) new infinite horizon discounted dynamic programming equation, the associated contractive property, and a new policy iteration algorithm. Finally, we provide illustrative examples for both the finite and infinite horizon cases. For the infinite horizon case we invoke the new policy iteration algorithm to compute the optimal strategies.

Key words. Stochastic Control, Minimax, Dynamic Programming, Total Variational Distance

AMS subject classifications. 90C39, 93E20, 49J35

1. Introduction. Dynamic programming recursions are often employed in optimal control and decision theory to establish existence of optimal strategies, to derive necessary and sufficient optimality conditions, and to compute the optimal strategies either in closed form or via algorithms [7, 15, 20]. The cost-to-go and the corresponding dynamic programming recursion, in their general form, are functionals of the conditional distribution of the underlying state process (controlled process) given the past and present state and control processes [7]. Thus, any ambiguity of the controlled process conditional distribution will affect the optimality of the strategies. The term "ambiguity" is used to differentiate from the term "uncertainty" often used in control nomenclature to account for situations in which the true and nominal distribution (induced by models) are absolutely continuous, and hence they are defined on the same state space. This distinction is often omitted from various robust deterministic and stochastic control approaches, including minimax and risk-sensitive formulations [1–4, 8, 9, 11, 13, 14, 16, 17, 19, 21]. In this paper, the class of models is described by a ball with respect to the total variation distance between the nominal distribution and the true distribution, hence it admits distributions which are singular with respect to the nominal distribution.

The main objective of this paper is to investigate the effect on the cost-to-go and dynamic programming of the ambiguity in the controlled process conditional distribution, and hence on the optimal decision strategies. Specifically, we quantify the conditional distribution ambiguity of the controlled process by a ball with respect to the total variation distance metric, centered at a nominal conditional distribution, and then we derive a new dynamic programming using minimax theory, with two players: player I the control process and player II the

*Department of Electrical and Computer Engineering, University of Cyprus (UCY), Nicosia, Cyprus. (tzortzis.ioannis@ucy.ac.cy).

†Department of Electrical and Computer Engineering, University of Cyprus (UCY), Nicosia, Cyprus. (chadcha@ucy.ac.cy).

‡School of Electrical Engineering, Royal Institute of Technology (KTH), Stockholm, Sweden. (themisc@kth.se).

conditional distribution (controlled process), opposing each other actions. In this minimax game formulation, player's I objective is to minimize the cost-to-go, while player's II objective is to maximize it. The maximization over the total variation distance ball of player II is addressed by first deriving results related to the maximization of linear functionals on a subset of the space of signed measures. Utilizing these results, a new dynamic programming recursion is presented which, in addition to the standard terms, includes additional terms that codify the level of ambiguity allowed by player II with respect to the total variation distance ball. Thus, the effect of player I, the control process, is to minimize, in addition to the classical terms, the difference between the maximum and minimum values of the cost-to-go, scaled by the radius of the total variation distance ambiguity set. We treat in a unified way the finite horizon case, under both the Markovian and non-Markovian nominal controlled processes, and the infinite horizon case. For the infinite horizon case we consider a discounted pay-off and we show that the operator associated with the resulting dynamic programming equation under total variation distance ambiguity is contractive. Consequently, we derive a new policy iteration algorithm to compute the optimal strategies. Finally, we provide examples for the finite and for the infinite horizon case.

Previous related work on optimization of stochastic systems subject to total variation distance ambiguity is found in [18] for continuous time controlled diffusion processes described by Itô differential equations. However, the solution method employed in [18] is fundamentally different; it approaches the maximization problem indirectly, by employing Large Deviations concepts to derive the maximizing measure as a convex combination of a tilted probability measure and the nominal measure, under restrictions on the class of measures considered. The dynamic programming equation derived in [18] is limited by the assumption that the maximizing measure is absolutely continuous with respect to the nominal measure.

In this paper, our focus is to understand the effect of total variation distance ambiguity of the conditional distribution on dynamic programming, from a different point of view, utilizing concepts from signed measures. Consequently, we derive a new dynamic programming recursion which depends explicitly on the radius of the total variation distance, the closed form expression of the maximizing measure, or the oscillator seminorm of the value function. One of the fundamental properties of the maximizing conditional distribution is that, as the ambiguity radius increases, the maximizing conditional distribution becomes singular with respect to the nominal distribution. The point to be made here is that the total variation distance ambiguity set admits controlled process distributions which are not necessarily defined on the same state space as the nominal controlled process distribution. In terms of robustness of the optimal policies, this additional feature is very attractive compared to minimax techniques based on relative entropy uncertainty or risk-sensitive pay-offs [1–4,8,9,11,13,14,16,17,19,21], because often the true controlled distribution lies on a higher dimensional state space compared to the nominal controlled process distribution.

The rest of the paper is organized as follows. In Section 1.1, we give a high level discussion on classical dynamic programming for MCM and we present some aspects of the problems and results obtained in the paper. In Section 2, we describe the abstract formulation of the minimax problem under total variation distance ambiguity, and we derive the closed form expression of the maximizing measure. In Section 3, we apply the abstract setup to Feedback Control Model (FCM) (e.g., non-Markov) and to MCM. We derive new dynamic programming recursions which characterize the optimality of minimax strategies. In Section 3.4, we treat the infinite horizon case, where we show that the dynamic programming operator is contractive, and we develop a new policy iteration algorithm. Finally, in Section 4 we present various examples to illustrate the applications of the new dynamic programming recursions.

1.1. Discussion on the Main Results. Next, we describe at a high level the results obtained in this paper.

1.1.1. Dynamic Programming of Finite Horizon Discounted-Markov Control Model. A finite horizon Discounted-Markov Control Model (D-MCM) with deterministic strategies is a septuple

$$\text{D-MCM}: \Big(\{\mathcal{X}_i\}_{i=0}^{n}, \{\mathcal{U}_i\}_{i=0}^{n-1}, \{\mathcal{U}_i(x_i): x_i \in \mathcal{X}_i\}_{i=0}^{n-1}, \{Q_i(dx_i|x_{i-1},u_{i-1}): (x_{i-1},u_{i-1}) \in \mathcal{X}_{i-1}\times\mathcal{U}_{i-1}\}_{i=0}^{n}, \{f_i\}_{i=0}^{n-1}, h_n, \alpha\Big) \tag{1.1}$$

consisting of

(a) State Space. A sequence of Polish spaces (complete separable metric spaces) $\{\mathcal{X}_i : i = 0,\ldots,n\}$, which model the state space of the controlled random process $\{x_j \in \mathcal{X}_j : j = 0,\ldots,n\}$.

(b) Control or Action Space. A sequence of Polish spaces $\{\mathcal{U}_i : i = 0,\ldots,n-1\}$, which model the control or action set of the control random process $\{u_j \in \mathcal{U}_j : j = 0,\ldots,n-1\}$.

(c) Feasible Controls or Actions. A family $\{\mathcal{U}_i(x_i) : x_i \in \mathcal{X}_i\}$ of non-empty measurable subsets $\mathcal{U}_i(x_i)$ of $\mathcal{U}_i$, where $\mathcal{U}_i(x_i)$ denotes the set of feasible controls or actions, when the controlled process is in state $x_i \in \mathcal{X}_i$, and the feasible state-actions pairs defined by $\mathbb{K}_i \triangleq \Big\{(x_i,u_i) : x_i \in \mathcal{X}_i, u_i \in \mathcal{U}_i(x_i)\Big\}$ are measurable subsets of $\mathcal{X}_i \times \mathcal{U}_i, i = 0,\ldots,n-1$.

(d) Controlled Process Distribution. A collection of conditional distributions or stochastic kernels $Q_i(dx_i|x_{i-1},u_{i-1})$ on $\mathcal{X}_i$ given $(x_{i-1},u_{i-1}) \in \mathbb{K}_{i-1} \subseteq \mathcal{X}_{i-1}\times\mathcal{U}_{i-1}, i = 0,\ldots,n$. The controlled process distribution is described by the sequence of transition probability distributions $\{Q_i(dx_i|x_{i-1},u_{i-1}) : (x_{i-1},u_{i-1}) \in \mathbb{K}_{i-1}, i = 0,\ldots,n\}$.

(e) Cost-Per-Stage. A collection of non-negative measurable functions $f_j : \mathbb{K}_j \to [0,\infty]$, called the cost-per-stage, such that $f_j(x,\cdot)$ does not take the value $+\infty$ for each $x \in \mathcal{X}_j, j = 0,\ldots,n-1$. The running pay-off functional is defined in terms of $\{f_j : j = 0,\ldots,n-1\}$.

(f) Terminal Cost. A bounded measurable non-negative function $h_n : \mathcal{X}_n \to [0,\infty)$ called the terminal cost. The pay-off functional at the last stage is defined in terms of h_n.

(g) Discounting Factor. A real number $\alpha \in (0,1)$ called the discounting factor.

The definition of D-MCM envisions applications of systems described by discrete-time dynamical state space models, which include random external inputs, since such models give rise to a collection of controlled processes distributions $\{Q_i(dx_i|x_{i-1},u_{i-1}):(x_{i-1},u_{i-1}) \in \mathbb{K}_{i-1}, i = 0,\ldots,n\}$. For any integer $j \geq 0$, define the product spaces by $\mathcal{X}_{0,j} \triangleq \times_{i=0}^{j}\mathcal{X}_i$ and $\mathcal{U}_{0,j-1} \triangleq \times_{i=0}^{j-1}\mathcal{U}_i$. Define the discounted sample pay-off by

$$F_{0,n}^{\alpha}(x_0,u_0,x_1,u_1,\ldots,x_{n-1},u_{n-1},x_n) \triangleq \sum_{j=0}^{n-1}\alpha^j f_j(x_j,u_j) + \alpha^n h_n(x_n). \tag{1.2}$$

The goal in Markov controlled optimization with deterministic strategies is to choose a control strategy or policy $g \triangleq \{g_j : j = 0,1,\ldots,n-1\}$, $g_j : \mathcal{X}_{0,j}\times\mathcal{U}_{0,j-1} \longrightarrow \mathcal{U}_j(x_j)$, $u_j^g = g_j(x_0^g,x_1^g,\ldots,x_j^g,u_0^g,u_1^g,\ldots,u_{j-1}^g)$, $j = 0,1,\ldots,n-1$ so as to minimize the pay-off functional

$$\begin{aligned}\mathbb{E}\Big\{\sum_{j=0}^{n-1}\alpha^j f_j(x_j^g,u_j^g) + \alpha^n h_n(x_n^g)\Big\} = \int_{\mathcal{X}_0\times\mathcal{X}_1\times\ldots\times\mathcal{X}_n} \\ F_{0,n}^{\alpha}\Big(x_0,u_0^g(x_0),x_1,u_1^g(x_0,x_1),\ldots,x_{n-1},u_{n-1}^g(x_0,x_1,\ldots,x_{n-1}),x_n) \\ Q_0(dx_0)Q_1(dx_1|x_0,u_0^g(x_0))\ldots Q_n(dx_n|x_{n-1},u_{n-1}^g(x_0,x_1,\ldots,x_{n-1}))\Big).\end{aligned} \tag{1.3}$$

Clearly, pay-off (1.3) is a functional of the collection of conditional distributions $\{Q_i(\cdot|\cdot) : i = 0, 1, \ldots, n\}$. Moreover, if this collection of distribution has countable support for each (x_{i-1}, u_{i-1}), $i = 0, \ldots, n$, then each integral in (1.3) is reduced to a countable summation.

For $(i, x) \in \{0, 1 \ldots, n\} \times \mathcal{X}_i$, let $V_i^0(x) \in \mathbb{R}$ represent the minimal cost-to-go or value function on the time horizon $\{i, i+1, \ldots, n\}$ if the controlled process starts at state $x_i = x$ at time i, defined by

$$V_i^0(x) \triangleq \inf_{\substack{g_k \in \mathcal{U}_k(x_k) \\ k=i,\ldots,n-1}} \mathbb{E}_{i,x}^g \Big\{ \sum_{j=i}^{n-1} \alpha^j f_j(x_j^g, u_j^g) + \alpha^n h_n(x_n^g) \Big\} \tag{1.4}$$

where $\mathbb{E}_{i,x}^g\{\cdot\}$ denotes expectation conditioned on $x_i^g = x$. A Markov property on the controlled process distributions, i.e., $Q_i(dx_i|x^{i-1}, u^{i-1})$=$Q_i(dx_i|x_{i-1}, u_{i-1})$, $\forall (x^{i-1}, u^{i-1}) \in \times_{j=0}^{i-1} \mathbb{K}_j$, $i = 0, 1, \ldots, n$, under admissible non-Markov strategies, implies that Markov control strategies are optimal [15]. Consequently, it can be shown that the value function (1.4) satisfies the following dynamic programming recursion relating the value functions $V_i^0(\cdot)$ and $V_{i+1}^0(\cdot)$ [15],

$$V_n^0(x) = \alpha^n h_n(x), \qquad x \in \mathcal{X}_n \tag{1.5}$$

$$V_i^0(x) = \inf_{u \in \mathcal{U}_i(x)} \Big\{ \alpha^i f_i(x, u) + \int_{\mathcal{X}_{i+1}} V_{i+1}^0(z) Q_{i+1}(dz|x, u) \Big\}, \qquad x \in \mathcal{X}_i. \tag{1.6}$$

Since the value function $V_i^0(x)$ defined by (1.4) and the dynamic programming recursion (1.5), (1.6) depend on the complete knowledge of the collection of conditional distributions $\{Q_i(\cdot|\cdot) : i = 0, \ldots, n\}$, any mismatch of the collection $\{Q_i(\cdot|\cdot) : i = 0, \ldots, n\}$ from the true collection of conditional distributions, will affect the optimality of the control strategies. Our objective is to address the impact of any ambiguity measured by the total variation distance between the true conditional distribution and a given nominal distribution on the cost-to-go (1.4), and dynamic programming recursion (1.5), (1.6).

1.1.2. Dynamic Programming of Infinite Horizon D-MCM. The infinite horizon D-MCM with deterministic strategies is a special case of the finite horizon D-MCM specified by a six-tuple

$$\Big(\mathcal{X}, \mathcal{U}, \{\mathcal{U}(x) : x \in \mathcal{X}\}, \{Q(dz|x, u) : (x, u) \in \mathcal{X} \times \mathcal{U}\}, f, \alpha\Big) \tag{1.7}$$

where the elements defined under (a)-(f) are independent of time index i. That is, the state space is $\mathcal{X}$, the control or action space is $\mathcal{U}$, the feasible controls or actions is a family $\{\mathcal{U}(x) : x \in \mathcal{X}\} \subset \mathcal{U}$, the controlled process distribution is a stochastic kernel $Q(\cdot|\cdot)$ on $\mathcal{X}$ given $\mathbb{K}$, where $\mathbb{K} \triangleq \Big\{(x, u) : x \in \mathcal{X}, u \in \mathcal{U}(x)\Big\}$, the cost-per-stage is a one stage cost $f : \mathbb{K} \longrightarrow [0, \infty]$, and there is no terminal cost (it is set to zero).

The dynamic programming equation of the infinite horizon D-MCM as given by [20] is a function $v_\infty^0 : \mathcal{X} \longrightarrow \mathbb{R}$ satisfying

$$v_\infty^0(x) = \inf_{u \in \mathcal{U}(x)} \Big\{ f(x, u) + \alpha \int_{\mathcal{X}} v_\infty^0(z) Q(dz|x, u) \Big\}, \qquad x \in \mathcal{X}. \tag{1.8}$$

Similarly to the finite horizon D-MCM, the dynamic programming equation (1.8) depends on the conditional distribution $Q(dz|x, u)$, hence any ambiguity or mismatch of $Q(dz|x, u)$ from the true distribution affects optimality of the strategies.

1.1.3. Dynamic Programming with Total Variation Distance Ambiguity. Motivated by the above discussion, the objective of this paper is to investigate dynamic programming under ambiguity of the conditional distributions of the controlled processes

$$\Big\{Q_i(dx_i|x_{i-1},u_{i-1}) : (x_{i-1},u_{i-1}) \in \mathbb{K}_{i-1}\Big\}, \quad i=0,\ldots,n.$$

The ambiguity of the conditional distributions of the controlled process is modeled by the total variation distance. Specifically, given a collection of nominal controlled process distributions $\{Q_i^o(dx_i|x_{i-1},u_{i-1}) : (x_{i-1},u_{i-1}) \in \mathbb{K}_{i-1}\}$, $i=0,\ldots,n$, the corresponding collection of true controlled process distributions $\{Q_i(dx_i|x_{i-1},u_{i-1}) : (x_{i-1},u_{i-1}) \in \mathbb{K}_{i-1}\}$, $i=0,\ldots,n$, is modeled by a set described by the total variation distance centered at the nominal conditional distribution having radius $R_i \in [0,2]$, $i=0,\ldots,n$, defined by

$$\mathbf{B}_{R_i}(Q_i^o)(x_{i-1},u_{i-1}) \triangleq \Big\{Q_i(\cdot|x_{i-1},u_{i-1}) : ||Q_i(\cdot|x_{i-1},u_{i-1}) - Q_i^o(\cdot|x_{i-1},u_{i-1})||_{TV} \leq R_i\Big\}.$$

Here $||\cdot||_{TV}$ denotes the total variation distance between two probability measures, $||\cdot||_{TV} : \mathcal{M}_1(\Sigma) \times \mathcal{M}_1(\Sigma) \longmapsto [0,\infty]$ defined by

$$||\alpha-\beta||_{TV} \triangleq \sup_{P\in\mathcal{P}(\Sigma)} \sum_{F_i\in P} |\alpha(F_i)-\beta(F_i)|, \quad \alpha,\beta \in \mathcal{M}_1(\Sigma) \tag{1.9}$$

where $\mathcal{M}_1(\Sigma)$ denotes the set of probability measures on $\mathcal{B}(\Sigma)$ and $\mathcal{P}(\Sigma)$ denotes the collection of all finite partitions of Σ. Note that the distance metric (1.9) induced by the total variation norm does not require absolute continuity of the measures $\alpha \in \mathcal{M}_1(\Sigma)$ and $\beta \in \mathcal{M}_1(\Sigma)$. The total variation distance model of ambiguity is quite general, and it includes linear, non-linear, finite and/or countable state space models, etc, since no assumptions are impossed on the structure of the stochastic control dynamical system model, which induces the collection of conditional distributions $\{Q_i(\cdot|\cdot) : i=0,\ldots,n\}$, $\{Q_i^o(\cdot|\cdot) : i=0,\ldots,n\}$. Given the above description of ambiguity in distribution, we re-formulate the value function and dynamic programming recursion via minimax theory as follows.

For $(i,x) \in \{0,1\ldots,n\} \times \mathcal{X}_i$, let $V_i(x) \in \mathbb{R}$ represent the minimal cost-to-go on the time horizon $\{i,i+1,\ldots,n\}$ if the state of the controlled process starts at state $x_i = x$ at time i, defined by

$$V_i(x) \triangleq \inf_{\substack{g_k\in\mathcal{U}_k(x_k)\\ k=i,\ldots,n-1}} \sup_{\substack{Q_{k+1}(\cdot|x_k,u_k)\in\mathbf{B}_{R_{k+1}}(Q_{k+1}^o)(x_k,u_k)\\ k=i,\ldots,n-1}} \mathbb{E}_{i,x}^g\Big\{\sum_{j=i}^{n-1}\alpha^j f_j(x_j^g,u_j^g) + \alpha^n h_n(x_n^g)\Big\}$$

where $\mathbb{E}_{i,x}^g$ denotes conditional expectation with respect to the true collection of conditional distribution $\{Q_k(\cdot|\cdot) : k=i,\ldots,n\}$. Even in the above minimax setting the Markov property of the controlled process distribution under an admissible non-Markov strategy implies that Markov control strategies are optimal. Moreover, the value function satisfies the following dynamic programming recursion relating the value function $V_i(\cdot)$ and $V_{i+1}(\cdot)$, for all $i=0,1,\ldots,n-1$.

$$V_n(x) = \alpha^n h_n(x), \quad x \in \mathcal{X}_n$$

$$V_i(x) = \inf_{u\in\mathcal{U}_i(x)} \sup_{Q_{i+1}(\cdot|x,u)\in\mathbf{B}_{R_{i+1}}(Q_{i+1}^o)(x,u)} \Big\{\alpha^i f_i(x,u) + \int_{\mathcal{X}_{i+1}} V_{i+1}(z)Q_{i+1}(dz|x,u)\Big\}, \; x \in \mathcal{X}_i.$$

Based on this formulation, if $V_{i+1}(\cdot)$ is bounded continuous non-negative, we show that the new dynamic programming equation is given by

$$V_n(x) = \alpha^n h_n(x), \quad x \in \mathcal{X}_n \tag{1.10}$$

$$\begin{aligned} V_i(x) = \inf_{u\in\mathcal{U}_i(x)} \Big\{ & \alpha^i f_i(x,u) + \int_{\mathcal{X}_{i+1}} V_{i+1}(z) Q^o_{i+1}(dz|x,u) \\ & + \frac{R_i}{2}\Big(\sup_{z\in\mathcal{X}_{i+1}} V_{i+1}(z) - \inf_{z\in\mathcal{X}_{i+1}} V_{i+1}(z)\Big)\Big\}, \quad x \in \mathcal{X}_i. \end{aligned} \tag{1.11}$$

Note that the new term in the right side of (1.11) is the oscillator seminorm of $V_{j+1}(\cdot)$ called the global modulus of continuity of $V_{j+1}(\cdot)$, which measures the difference between the maximum and minimum values of $V_{j+1}(\cdot)$.

For the infinite horizon D-MCM the new dynamic programming equation is given by

$$\begin{aligned} v_\infty(x) = \inf_{u\in\mathcal{U}(x)} \Big\{ & f(x,u) \\ & + \alpha \int_{\mathcal{X}} v_\infty(z) Q^o(dz|x,u) + \alpha \frac{R}{2}\Big(\sup_{z\in\mathcal{X}} v_\infty(z) - \inf_{z\in\mathcal{X}} v_\infty(z)\Big)\Big\}, \quad x \in \mathcal{X}. \end{aligned} \tag{1.12}$$

For finite and countable alphabet spaces $\mathcal{X}_j$, $\mathcal{X}$, the integrals in the right hand side of (1.11), (1.12) are replaced by summations.

In addition to the D-MCM, we will also discuss the general discounted feedback control model (i.e., we relax the Markovian assumption). In summary, the issues discussed and results obtained in this paper are the following: (1) formulation of finite horizon discounted stochastic optimal control subject to conditional distribution ambiguity described by total variation distance via minimax theory; (2) dynamic programming recursions for a) nominal D-MCM, and b) Discounted-Feedback Control Model (D-FCM), under total variation distance ambiguity on the conditional distribution of the controlled process; (3) formulation of the infinite horizon D-MCM and dynamic programming equation under conditional distribution ambiguity described by total variation distance via minimax theory; (4) characterization of the maximizing conditional distribution belonging to the total variation distance set, and the corresponding new dynamic programming recursions; (5) contraction property of the infinite horizon D-MCM dynamic programming and new policy iteration algorithm; (6) examples for the finite and infinite horizon cases.

2. Maximization With Total Variation Distance Ambiguity. In this section, we recall certain results from [18] on the maximization of a linear functional on the space of probability distributions subject to total variation distance ambiguity. We use these results to derive the maximizing probability distribution subject to total variation distance ambiguity of the controlled process.

Let (Σ, d_Σ) denote a complete, separable metric space (a Polish space), and $(\Sigma, \mathcal{B}(\Sigma))$ the corresponding measurable space, in which $\mathcal{B}(\Sigma)$ is the σ-algebra generated by open sets in Σ. Let $\mathcal{M}_1(\Sigma)$ denote space of countably additive probability measures on $(\Sigma, \mathcal{B}(\Sigma))$. Define the spaces

$$BC(\Sigma) \triangleq \big\{\text{Bounded continuous functions } \ell : \Sigma \longrightarrow \mathbb{R} : ||\ell|| \triangleq \sup_{x\in\Sigma} |\ell(x)| < \infty\big\}$$

$$BM(\Sigma) \triangleq \big\{\text{Bounded measurable functions } \ell : \Sigma \longrightarrow \mathbb{R} : ||\ell|| < \infty\big\}$$

$$C(\Sigma) \triangleq \big\{\text{Continuous functions } \ell : \Sigma \longrightarrow \mathbb{R} : ||\ell|| < \infty\big\}, \quad C^+(\Sigma) \triangleq \big\{\ell \in C(\Sigma) : \ell \geq 0\big\}$$

$$BC^+(\Sigma) \triangleq \big\{\ell \in BC(\Sigma) : \ell \geq 0\Big\}, \quad BM^+(\Sigma) \triangleq \big\{\ell \in BM(\Sigma) : \ell \geq 0\big\}.$$

Clearly, $BC(\Sigma)$, $BM(\Sigma)$, $C(\Sigma)$ are Banach spaces. We present the maximizing measure for $\ell \in BC^+(\Sigma)$, although the results can be generalized to real-valued functions $\ell \in L^{\infty,+}(\Sigma, \mathcal{B}(\Sigma), \nu)$, the set of all $\mathcal{B}(\Sigma)$-measurable, non-negative essentially bounded functions defined $\nu-a.e.$ endowed with the essential supremum norm $||\ell||_{\infty,\nu}=\nu$-ess $\sup_{x\in\Sigma} \ell(x)$.

From [18], we have the following. For $\ell \in BC^+(\Sigma)$, and $\mu \in \mathcal{M}_1(\Sigma)$ fixed, then

$$L(\nu^*) \triangleq \sup_{||\nu-\mu||_{TV}\leq R} \int_\Sigma \ell(x)\nu(dx) = \frac{R}{2}\Big\{ \sup_{x\in\Sigma} \ell(x) - \inf_{x\in\Sigma} \ell(x)\Big\} + \int_\Sigma \ell(x)\mu(dx) \tag{2.1}$$

where $R \in [0, 2]$, ν^* satisfies the constraint $||\xi^*||_{TV} = ||\nu^* - \mu||_{TV} = R$, it is normalized $\nu^*(\Sigma) = 1$, and $\nu^*(A) \in [0, 1]$ on any $A \in \mathcal{B}(\Sigma)$. Moreover, by defining [1]

$$x^0 \in \Sigma^0 \triangleq \{x \in \Sigma : \ell(x) = \sup\{\ell(x) : x \in \Sigma\} \equiv \ell_{\max}\}$$
$$x_0 \in \Sigma_0 \triangleq \{x \in \Sigma : \ell(x) = \inf\{\ell(x) : x \in \Sigma\} \equiv \ell_{\min}\}$$

then, the pay-off $L(\nu^*)$ can be written as

$$L(\nu^*) = \int_{\Sigma^0} \ell_{\max}\nu^*(dx) + \int_{\Sigma_0} \ell_{\min}\nu^*(dx) + \int_{\Sigma\setminus\Sigma^0\cup\Sigma_0} \ell(x)\mu(dx) \tag{2.2}$$

and the optimal distribution $\nu^* \in \mathcal{M}_1(\Sigma)$, which satisfy the total variation constraint, is given by

$$\begin{gathered}\int_{\Sigma^0} \nu^*(dx) = \mu(\Sigma^0) + \frac{R}{2} \in [0, 1], \qquad \int_{\Sigma_0} \nu^*(dx) = \mu(\Sigma_0) - \frac{R}{2} \in [0, 1] \\ \nu^*(A) = \mu(A), \quad \forall A \subseteq \Sigma \setminus \Sigma^0 \cup \Sigma_0.\end{gathered} \tag{2.3}$$

Note that if $\Sigma^0 = \Sigma_0 = \{\emptyset\}$ then $\nu(\Sigma^0) = \nu(\Sigma_0) = 0$, and $L(\nu^*) = \int_{\Sigma\setminus\Sigma^0\cup\Sigma_0} \ell(x)\mu(dx)$.

The second right hand side term in (2.1) is related to the oscillator semi-norm of $f \in BM(\Sigma)$, called the global modulus of continuity, and it is defined by

$$\text{osc}(f) \triangleq \sup_{(x,y)\in\Sigma\times\Sigma} |f(x) - f(y)| = 2 \inf_{\beta\in\mathbb{R}} ||f - \beta||, \qquad \text{for } f \in BM(\Sigma).$$

However, for $f \in BM^+(\Sigma)$ then

$$\text{osc}(f) = \sup_{x\in\Sigma} |f(x)| - \inf_{x\in\Sigma} |f(x)| = \sup_{x\in\Sigma} f(x) - \inf_{x\in\Sigma} f(x).$$

Note that the above results can be extended to $f \in C^+(\Sigma)$.

The Maximizing Measure for Finite and Countable Alphabet Spaces

Here, we further elaborate on the form of the maximizing measures for finite and countable alphabet spaces, since we use them to analyze finite horizon D-MCM and D-FCM, and infinite horizon D-MCM with finite (or countable) state and control spaces.

Let Σ be a non-empty denumerable set endowed with the discrete topology including finite cardinality $|\Sigma|$, with $\mathcal{M}_1(\Sigma)$ identified with the standard probability simplex in $\mathbb{R}^{|\Sigma|}$. That is, the set of all $|\Sigma|$-dimensional vectors which are probability vectors, $\{\nu(x) : x \in \Sigma\} \in \mathcal{M}_1(\Sigma)$, $\{\mu(x) : x \in \Sigma\} \in \mathcal{M}_1(\Sigma)$, and let $\ell \triangleq \{\ell(x) : x \in \Sigma\} \in \mathbb{R}_+^{|\Sigma|}$. Define the maximum and minimum values of $\{\ell(x) : x \in \Sigma\}$ by

$$\ell_{\max} \triangleq \max_{x\in\Sigma} \ell(x), \quad \ell_{\min} \triangleq \min_{x\in\Sigma} \ell(x)$$

[1] We adopt the standard definitions; infimum (supremum) of an empty set to be $+\infty$ ($-\infty$).

and its corresponding support sets by

$$\Sigma^0 \triangleq \big\{x \in \Sigma : \ell(x) = \ell_{\max}\big\}, \quad \Sigma_0 \triangleq \big\{x \in \Sigma : \ell(x) = \ell_{\min}\big\}.$$

For all remaining sequence, $\big\{\ell(x) : x \in \Sigma \setminus \Sigma^0 \cup \Sigma_0\big\}$, and for $1 \leq r \leq |\Sigma \setminus \Sigma^0 \cup \Sigma_0|$, define recursively the set of indices for which the sequence achieves its $(k+1)^{th}$ smallest value by

$$\Sigma_k \triangleq \Big\{x \in \Sigma : \ell(x) = \min\big\{\ell(\alpha) : \alpha \in \Sigma \setminus \Sigma^0 \cup (\bigcup_{j=1}^{k} \Sigma_{j-1})\big\}\Big\}, \quad k \in \{1, 2, \ldots, r\}$$

till all the elements of Σ are exhausted. Further, define the corresponding values of the sequence on sets Σ_k by

$$\ell(\Sigma_k) \triangleq \min_{x \in \Sigma \setminus \Sigma^0 \cup (\bigcup_{j=1}^{k} \Sigma_{j-1})} \ell(x), \qquad k \in \{1, 2, \ldots, r\}$$

where r is the number of Σ_k sets which is at most $|\Sigma \setminus \Sigma^0 \cup \Sigma_0|$. For example, when $k = 1$, $\ell(\Sigma_1) = \min_{x \in \Sigma \setminus \Sigma^0 \cup \Sigma_0} \ell(x)$, when $k = 2$, $\ell(\Sigma_2) = \min_{x \in \Sigma \setminus \Sigma^0 \cup \Sigma_0 \cup \Sigma_1} \ell(x)$ and so on.

In [10] it is shown that the maximum pay-off subject to total variation constraint is given by

$$L(\nu^*) = \ell_{\max}\nu^*(\Sigma^0) + \ell_{\min}\nu^*(\Sigma_0) + \sum_{k=1}^{r} \ell(\Sigma_k)\nu^*(\Sigma_k), \tag{2.4}$$

and that the optimal probabilities are given by (a water-filling) the following equations.

$$\nu^*(\Sigma^0) \triangleq \sum_{x \in \Sigma^0} \nu^*(x) = \sum_{x \in \Sigma^0} \mu(x) + \frac{\alpha}{2} \equiv \mu(\Sigma^0) + \frac{\alpha}{2} \tag{2.5}$$

$$\nu^*(\Sigma_0) \triangleq \sum_{x \in \Sigma_0} \nu^*(x) = \Big(\sum_{x \in \Sigma_0} \mu(x) - \frac{\alpha}{2}\Big)^+ \equiv \Big(\mu(\Sigma_0) - \frac{\alpha}{2}\Big)^+ \tag{2.6}$$

$$\begin{aligned} \nu^*(\Sigma_k) \triangleq \sum_{x \in \Sigma_k} \nu^*(x) &= \Big(\sum_{x \in \Sigma_k} \mu(x) - \Big(\frac{\alpha}{2} - \sum_{j=1}^{k} \sum_{x \in \Sigma_{j-1}} \mu(x)\Big)^+\Big)^+ \\ &\equiv \Big(\mu(\Sigma_k) - \Big(\frac{\alpha}{2} - \sum_{j=1}^{k} \mu(\Sigma_{j-1})\Big)^+\Big)^+ \end{aligned} \tag{2.7}$$

$$\alpha \triangleq \min(R, R_{\max}), \;\; R_{\max} \triangleq 2(1 - \sum_{x \in \Sigma^0} \mu(x)) \equiv 2(1 - \mu(\Sigma^0)), \quad R \in [0, 2] \tag{2.8}$$

where $k \in \{1, 2, \ldots, r\}$ and r is the number of Σ_k sets which is at most $|\Sigma \setminus \Sigma^0 \cup \Sigma_0|$.

The parameter α reinforces the intuitive notion of the total variation between the true and nominal probability distribution as having attributes similar to "physical mass". Thus, if $\alpha = R_{\max}$, then (2.5) implies that the probability "mass" on Σ^0 set is $\nu^*(\Sigma^0) = 1$ and hence $\nu^*(\Sigma \setminus \Sigma^0) = 0$. However, if $\alpha = R < R_{\max}$, then (2.5) implies that the probability "mass" on Σ^0 set is $\nu^*(\Sigma^0) < 1$ and hence equations (2.6)-(2.7) are employed. While $\nu^*(\Sigma_0) > 0$, (2.7) implies that $\nu^*(\Sigma_k) = \mu(\Sigma_k)$ for all $k = 1, \ldots, r$. However, if $\nu^*(\Sigma_0) = 0$, that is, all the probability "mass" is removed from Σ_0, then the solution is obtained by moving further into the partition using (2.7). For all $R \in [0, 2]$, the resulting solution is described via a water-filling effect.

We are now equipped with the solution of maximizing linear functionals with total variation distance ambiguity for both finite, countable alphabets, and abstract alphabet spaces (Polish spaces), and therefore we are ready to apply these results to the dynamic programming recursion under ambiguity on the conditional distribution.

3. Minimax Stochastic Control with Total Variation Distance Ambiguity. In this section, we first introduce the general definition of finite horizon Discounted-Feedback Control Model (D-FCM) with randomized and deterministic control policies, under total variation distance uncertainty (which includes the D-MCM introduced in Section 1.1), and then we apply the characterization of the maximizing distribution of Section 2 to the dynamic programming recursion. In the last section we discuss the infinite horizon D-MCM.

Define $\mathbb{N}^n \triangleq \{0,1,2,\ldots,n\}, n \in \mathbb{N}$. The state space and the control space are sequences of Polish spaces $\{\mathcal{X}_j : j=0,1,\ldots,n\}$ and $\{\mathcal{U}_j : j=0,1,\ldots,n-1\}$, respectively. These spaces are associated with their corresponding measurable spaces $(\mathcal{X}_j, \mathcal{B}(\mathcal{X}_j)), \forall j \in \mathbb{N}^n$, $(\mathcal{U}_j, \mathcal{B}(\mathcal{U}_j)), \forall j \in \mathbb{N}^{n-1}$. Define the product spaces by $\mathcal{X}_{0,n} \triangleq \times_{i=0}^{n} \mathcal{X}_i, \mathcal{U}_{0,n-1} \triangleq \times_{i=0}^{n-1} \mathcal{U}_i$, and introduce their product measurable spaces, $(\mathcal{X}_{0,n}, \mathcal{B}(\mathcal{X}_{0,n}))$, $(\mathcal{U}_{0,n-1}, \mathcal{B}(\mathcal{U}_{0,n-1}))$, respectively, for $n \in \mathbb{N}^n$. The state process is denoted by $x^n \triangleq \{x_j : j=0,1,\ldots,n\}$, and the control process is denoted by $u^{n-1} \triangleq \{u_j : j=0,1,\ldots,n-1\}$. For any measurable spaces $(\mathcal{X}, \mathcal{B}(\mathcal{X})), (\mathcal{Y}, \mathcal{B}(\mathcal{Y}))$, the set of stochastic Kernels on $(\mathcal{Y}, \mathcal{B}(\mathcal{Y}))$ conditioned on $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$ is denoted by $\mathcal{Q}(\mathcal{Y}|\mathcal{X})$.

Given $(\mathcal{X}_{0,n}, \mathcal{B}(\mathcal{X}_{0,n}))$, $(\mathcal{U}_{0,n-1}, \mathcal{B}(\mathcal{U}_{0,n-1}))$ the Borel state and control or action spaces, respectively, and the initial state distribution $\nu_0(dx_0)$, we introduce the space $H_{0,n}$ of admissible observable histories by

$$H_{0,n} \triangleq \mathbb{K}_0 \times \mathbb{K}_1 \times \ldots \times \mathbb{K}_{n-1} \times \mathcal{X}_n \equiv \times_{i=0}^{n-1} \mathbb{K}_i \times \mathcal{X}_n, \quad n \in \mathbb{N}, \quad H_{0,0} = \mathcal{X}_0$$

where $\mathbb{K}_i \triangleq \{(x_i, u_i) : x_i \in \mathcal{X}_i, u_i \in \mathcal{U}_i(x_i)\}$, denote the feasible state-action pairs, for $i = 0,1,\ldots,n-1$. A typical element $h_{0,n} \in H_{0,n}$ is a sequence of the form

$$h_{0,n} = (x_0, u_0, \ldots, x_{n-1}, u_{n-1}, x_n), \quad (x_i, u_i) \in \mathbb{K}_i, \quad i = 0,\ldots,n-1, \quad x_n \in \mathcal{X}_n.$$

Similarly, introduce

$$G_{0,n} = \mathcal{X}_0 \times \mathcal{U}_0 \times \ldots \times \mathcal{X}_{n-1} \times \mathcal{U}_{n-1} \times \mathcal{X}_n \equiv \times_{i=0}^{n-1} (\mathcal{X}_i \times \mathcal{U}_i) \times \mathcal{X}_n, \quad n \in \mathbb{N}$$
$$G_{0,0} = H_{0,0} = \mathcal{X}_0.$$

The spaces $G_{0,n}$ and $H_{0,n}$ are equipped with the natural σ-algebra $\mathcal{B}(G_{0,n})$ and $\mathcal{B}(H_{0,n})$, respectively.

Next, we give the precise definition of discounted feedback control model.

DEFINITION 3.1. *A finite horizon D-FCM is a septuple*

$$\textit{D-FCM}: \Big(\mathcal{X}_{0,n}, \mathcal{U}_{0,n-1}, \{\mathcal{U}_i(x_i) : x_i \in \mathcal{X}_i\}_{i=0}^{n-1}, \{Q_i(dx_i|x^{i-1}, u^{i-1}) : (x^{i-1}, u^{i-1}) \in \mathcal{X}_{0,i-1} \times \mathcal{U}_{0,i-1}\}_{i=0}^{n}, \{f_i\}_{i=0}^{n-1}, h_n, \alpha\Big) \tag{3.1}$$

*consisting of the items **(a)-(c)**, **(e)-(g)** of finite horizon D-MCM (1.1), while the controlled process distribution in **(d)** is replaced by the non-Markov collection* $\{Q_i(dx_i|x^{i-1}, u^{i-1}) : (x^{i-1}, u^{i-1}) \in \times_{j=0}^{i-1} \mathbb{K}_j\}_{i=0}^{n}$.

Next, we give the definitions of randomized, deterministic, and stationary control strategies or policies.

DEFINITION 3.2. *A randomized control strategy is a sequence* $\pi \triangleq \{\pi_0, \ldots, \pi_{n-1}\}$ *of stochastic kernels* $\pi_i(\cdot|\cdot)$ *on* $(\mathcal{U}_i, \mathcal{B}(\mathcal{U}_i))$ *conditioned on* $(H_{0,i}, \mathcal{B}(H_{0,i}))$ *(e.g.,* $\pi_i(du_i|x^i, u^{i-1})$ *) satisfying*

$$\pi_i(\mathcal{U}_i(x_i)|x^i, u^{i-1}) = 1 \quad \textit{for every} \quad (x^i, u^{i-1}) \in H_{0,i}, \quad i = 0,1,\ldots,n-1.$$

The set of all such policies is denoted by $\mathbf{\Pi}_{0,n-1}$.

A strategy $\pi \triangleq \{\pi_i : i = 0, \ldots, n-1\} \in \mathbf{\Pi}_{0,n-1}$ *is called*

(a) *deterministic feedback strategy if there exists a sequence* $g \triangleq \{g_j : j = 0, 1, \ldots, n-1\}$ *of measurable functions* $g_j : \times_{i=0}^{j-1} \mathbb{K}_i \times \mathcal{X}_j \longmapsto \mathcal{U}_j$, *such that for all* $\left(x^j, u^{j-1}\right) \in H_{0,j}$, $j \in \mathbb{N}^{n-1}$, $g_j(x_0, u_0, x_1, u_1, \ldots, x_{j-1}, u_{j-1}, x_j) \in \mathcal{U}_j(x_j)$, *and* $\pi_j\left(\cdot|x^j, u^{j-1}\right)$ *assigns mass 1 to some point in* $\mathcal{U}_j$, *that is,*

$$\pi_i\left(A_i|x^i, u^{i-1}\right) = I_{A_i}\left(g_i\left(x^i, u^{i-1}\right)\right), \quad \forall A_i \in \mathcal{B}(\mathcal{U}_i), \quad i = 0, 1, \ldots, n-1,$$

where $I_{A_i}(\cdot)$ *is the indicator function of* $A_i \in \mathcal{B}(\mathcal{U}_i)$.
The set of deterministic feedback strategies is denoted by $\mathbf{\Pi}^{DF}_{0,n-1}$;

(b) *deterministic Markov strategy if there exists a sequence* $g \triangleq \{g_j : j = 0, 1, \ldots, n-1\}$ *of measurable functions* $g_j : \mathcal{X}_j \to \mathcal{U}_j$ *satisfying* $g_j(x_j) \in \mathcal{U}_j(x_j)$ *for all* $x_j \in \mathcal{X}_j$, $j \in \mathbb{N}^{n-1}$, *and* $\pi_j(\cdot|x^j, u^{j-1})$ *is concentrated at* $g_j(x_j) \in \mathcal{U}_j(x_j)$ *for all* $(x^j, u^{j-1}) \in H_{0,j}$, $j \in \mathbb{N}^{n-1}$.
The set of deterministic Markov strategies is denoted by $\mathbf{\Pi}^{DM}_{0,n-1}$;

(c) *deterministic stationary Markov strategy if there exists a measurable function* $g : \mathcal{X} \longrightarrow \mathcal{U}$ *such that* $g(x_t) \in \mathcal{U}(x_t)$, $\forall x_t \in \mathcal{X}$, *and* $\pi_j(\cdot|x^j, u^{j-1})$ *assigns mass to some point* u_j, $\forall (x^j, u^{j-1}) \in H_{0,j}$, *e.g.,*

$$\pi_i(A_i|x^i, u^{i-1}) = I_{A_i}(g(x_i)), \quad \forall A_i \in \mathcal{B}(\mathcal{U}_i), \quad i = 0, \ldots, n-1.$$

The set of deterministic stationary Markov strategies is denoted by $\mathbf{\Pi}^{DS}_{0,n-1}$.

According to Definition 3.2, the set of control policies is non-empty, since we have assumed existence of measurable functions $g_j : \mathbb{K}_{0,j-1} \times \mathcal{X}_j \longrightarrow \mathcal{U}_j$ such that $\forall x^j, u^{j-1} \in \mathbb{K}_{0,j-1} \times \mathcal{X}_j$, $g_j(x^j, u^{j-1}) \in \mathcal{U}_j(\mathcal{X}_j), \forall j \in \mathbb{N}^{n-1}$. Sufficient conditions for this to hold are in general obtained via measurable selection theorems [12]. For denumerable set (countable alphabet) $\mathcal{X}_j$ endowed with the discrete topology any function is measurable. Given a controlled process $\left\{Q_i(\cdot|x^{i-1}, u^{i-1}) : (x^{i-1}, u^{i-1}) \in \mathbb{K}_{0,i-1}\right\}_{i=0}^{n}$ and a randomized control process $\left\{\pi_i(\cdot|x^i, u^{i-1}) : (x^i, u^{i-1}) \in \mathbb{K}_{0,i-1} \times \mathcal{X}_i\right\}_{i=0}^{n} \in \mathbf{\Pi}_{0,n-1}$ and the initial probability $\nu_0(\cdot) \in \mathcal{M}_1(\mathcal{X}_0)$, then by Ionescu-Tulceu theorem [6] there exists a unique probability measure $\mathbf{Q}^{\pi}_{\nu}$ on $(\Omega, \mathcal{F})$ defined by

$$\begin{aligned} \mathbf{Q}^{\pi}_{\nu}(dx_0, du_0, dx_1, du_1, \ldots, dx_{n-1}, du_{n-1}, dx_n) &= Q_0(dx_0)\pi_0(du_0|x_0) \\ \otimes Q_1(dx_1|x_0, u_0)\pi_1(du_1|x^1, u_0) \otimes \ldots \otimes Q_{n-1}(dx_{n-1}|x^{n-2}, u^{n-2}) & \\ \pi_{n-1}(du_{n-1}|x^{n-1}, u^{n-2}) \otimes Q_n(dx_n|x^{n-1}, u^{n-1}) & \end{aligned} \tag{3.2}$$

such that

$$\begin{aligned} &\mathbf{Q}^{\pi}_{\nu}(x_0 \in A) = \nu(A), \quad A \in \mathcal{B}(\mathcal{X}_0) \\ &\mathbf{Q}^{\pi}_{\nu}(u_j \in B|h_{0,j}) = \pi_t(B|h_{0,j}), \quad B \in \mathcal{B}(\mathcal{U}_j) \\ &\mathbf{Q}^{\pi}_{\nu}(x_{j+1} \in C|h_{0,j}, u_t) = Q(C|h_{0,j}, u_j), \quad C \in \mathcal{B}(\mathcal{X}_{j+1}). \end{aligned}$$

Given the sample pay-off

$$F^{\alpha}_{0,n}(x_0, u_0, x_1, u_1, \ldots, x_{n-1}, u_{n-1}, x_n) \triangleq \sum_{j=0}^{n-1} \alpha^j f_j(x_j, u_j) + \alpha^n h_n(x_n) \tag{3.3}$$

its expectation is

$$\begin{aligned} \mathbb{E}_{\mathbf{Q}^{\pi}_{\nu}}\left\{F^{\alpha}_{0,n}(x_0, u_0, x_1, u_1, \ldots, x_{n-1}, u_{n-1}, x_n)\right\} = \int F^{\alpha}_{0,n}(x_0, u_0, x_1, u_1, \ldots, & \\ x_{n-1}, u_{n-1}, x_n)\mathbf{Q}^{\pi}_{\nu}(dx_0, du_0, dx_1, du_1, \ldots, dx_{n-1}, du_{n-1}, dx_n). & \end{aligned} \tag{3.4}$$

Note that the class of randomized strategies $\mathbf{\Pi}_{0,n-1}$ embeds deterministic feedback and Markov strategies.

3.1. Variation Distance Ambiguity. Next, we introduce the definitions of nominal controlled process distributions (for finite horizon D-FCM and D-MCM), and their corresponding ambiguous controlled process distributions.

For each $\pi \in \mathbf{\Pi}^{DF}_{0,n-1}$, $\pi \in \mathbf{\Pi}^{DM}_{0,n-1}$ and $\pi \in \mathbf{\Pi}^{DS}_{0,n-1}$ the nominal controlled process is described by a sequence of conditional distributions as follows.

DEFINITION 3.3. *(Nominal Controlled Process Distributions). A nominal controlled state processes* $\{x^g = x^g_0, x^g_1, \ldots, x^g_n : \pi \in \mathbf{\Pi}^{DF}_{0,n-1}, \quad \pi \in \mathbf{\Pi}^{DM}_{0,n-1}, \quad or \quad \pi \in \mathbf{\Pi}^{DS}_{0,n-1}\}$ *corresponds to a sequence of stochastic kernels as follows:*
(a) *Feedback Controlled Process.*

$$\textit{For every} \quad A \in \mathcal{B}(\mathcal{X}_j), \quad Prob(x_j \in A | x^{j-1}, u^{j-1}) = Q^o_j(A | x^{j-1}, u^{j-1})$$

where $Q^o_j(A|x^{j-1}, u^{j-1}) \in \mathcal{Q}(\mathcal{X}_j|\mathbb{K}_{0,j-1}), \forall j \in \mathbb{N}^n_+$.
(b) *Markov Controlled Process.*

$$\textit{For every} \quad A \in \mathcal{B}(\mathcal{X}_j), \quad Prob(x_j \in A | x^{j-1}, u^{j-1}) = Q^o_j(A | x_{j-1}, u_{j-1})$$

where $Q^o_j(A|x_{j-1}, u_{j-1}) \in \mathcal{Q}(\mathcal{X}_j|\mathbb{K}_{j-1}), \forall j \in \mathbb{N}^n_+$.
(c) *Stationary Markov Controlled Process.*

$$\textit{For every} \quad A \in \mathcal{B}(\mathcal{X}), \quad Prob(x_j \in A | x^{j-1}, u^{j-1}) = Q^o(A | x_{j-1}, u_{j-1})$$

where $Q^o(A|x_{j-1}, u_{j-1}) \in \mathcal{Q}(\mathcal{X}|\mathbb{K})$.

The class of controlled processes is described by the sequence of stochastic kernels,

$$\{Q_j(dx_j|x^{j-1}, u^{j-1}) \in \mathcal{Q}(\mathcal{X}_j|\mathbb{K}_{0,j-1} : j = 0, \ldots, n\}$$

belonging to a total variation distance set as follows.

DEFINITION 3.4. *(Class of Controlled Process Distribution) Given a nominal controlled process stochastic kernel of Definition 3.3, and* $R_i \in [0, 2], 0 \leq i \leq n$ *the class of controlled process stochastic kernels is defined as follows:*
(a) *Class with respect to Feedback Nominal Controlled Process.*
Given a fixed $Q^o_j(\cdot|x^{j-1}, u^{j-1}) \in \mathcal{Q}(\mathcal{X}_j|\mathbb{K}_{0,j-1})$, $j = 0, 1, \ldots, n$ *the class of stochastic kernels is defined by*

$$\mathbf{B}_{R_i}(Q^o_i)(x^{i-1}, u^{i-1}) \triangleq \Big\{ Q_i(\cdot|x^{i-1}, u^{i-1}) \in \mathcal{Q}(\mathcal{X}_i|\mathbb{K}_{0,i-1}) : \\ ||Q_i(\cdot|x^{i-1}, u^{i-1}) - Q^o_i(\cdot|x^{i-1}, u^{i-1})||_{TV} \leq R_i \Big\}, \quad i = 0, 1, \ldots, n.$$

(b) *Class with respect to Markov Nominal Controlled Process.*
Given a fixed $Q^o_j(\cdot|x_{j-1}, u_{j-1}) \in \mathcal{Q}(\mathcal{X}_j|\mathbb{K}_{j-1})$, $j = 0, 1, \ldots, n$ *the class of stochastic kernels is defined by*

$$\mathbf{B}_{R_i}(Q^o_i)(x^{i-1}, u^{i-1}) \triangleq \Big\{ Q_i(\cdot|x^{i-1}, u^{i-1}) \in \mathcal{Q}(\mathcal{X}_i|\mathbb{K}_{0,i-1}) : \\ ||Q_i(\cdot|x^{i-1}, u^{i-1}) - Q^o_i(\cdot|x_{i-1}, u_{i-1})||_{TV} \leq R_i \Big\}, \quad i = 0, 1, \ldots, n.$$

***(c)** Class with respect to Stationary Markov Nominal Controlled Process.*
Given a fixed $Q^o(\cdot|x_{j-1},u_{j-1}) \in \mathcal{Q}(\mathcal{X},\mathbb{K})$ the class of stochastic kernels is defined by

$$\mathbf{B}_R(Q^0)(x,u) \triangleq \Big\{ Q(\cdot|x,u) \in \mathcal{Q}(\mathcal{X}|\mathbb{K}) : ||Q(\cdot|x,u) - Q^o(\cdot|x,u)||_{TV} \leq R \Big\}.$$

Note that in Definition 3.4 (a), (b), although we use the same notation $\mathbf{B}_{R_i}(Q_i^o)(x^{i-1},u^{i-1})$ these sets are different because the nominal distribution $Q_i^o(\cdot|\cdot)$ can be of Feedback or Markov form. The above model is motivated by the fact that dynamic programming involves conditional expectation with respect to the collection of conditional distributions $\{Q_i(\cdot|x^{i-1},u^{i-1}) \in \mathcal{Q}(\mathcal{X}_i|\mathbb{K}_{0,i-1}) : i=0,\ldots,n\}$. Therefore, any ambiguity in these distributions will affect the optimality of the strategies.

3.2. Pay-Off Functional. For each $\pi \in \mathbf{\Pi}_{0,n-1}^{DF}$ or $\pi \in \mathbf{\Pi}_{0,n-1}^{DM}$ the average pay-off is defined by

$$J_{0,n}(\pi, Q_i : i=0,\ldots,n) \triangleq \mathbb{E}_{\mathbf{Q}_\nu^\pi}\Big\{ \sum_{j=0}^{n-1} \alpha^j f_j(x_j,u_j) + \alpha^n h_n(x_n) \Big\} \tag{3.5}$$

where $\mathbb{E}_{\mathbf{Q}_\nu^\pi}\{\cdot\}$ denotes expectation with respect to the true joint measure $\mathbf{Q}_\nu^\pi(dx^n, du^{n-1})$ defined by (3.2) such that $Q_i(\cdot|x^{i-1},u^{i-1}) \in \mathbf{B}_{R_i}(Q_i^o)$, $i=0,1,\ldots,n$ (e.g., it belongs to the total variation distance ball of Definition 3.4).

Next, we introduce assumptions so that the maximization over the class of ambiguous measures is well-defined.

ASSUMPTION 3.5. *The nominal system family satisfies the following assumption: The maps $\{f_j : \mathcal{X}_j \times \mathcal{U}_j \longmapsto \mathbb{R} : j=0,1,\ldots,n-1\}$, $h_n : \mathcal{X}_n \longmapsto \mathbb{R}$ are bounded, continuous and non-negative.*

Note that it is possible to relax Assumption 3.5 to lower semi-continuous non-negative functions bounded from below.

3.3. Minimax Dynamic Programming for Finite Horizon D-FCM and D-MCM. In this section we shall apply the results of Section 2 to formulate and solve minimax stochastic control under a) finite horizon D-FCM ambiguity, and b) finite horizon D-MCM ambiguity.

3.3.1. Dynamic Programming for Finite Horizon D-FCM Subject to Ambiguity. Utilizing the above formulation, next we define the minimax stochastic control problem, where the maximization is over a total variation distance ball, centered at the nominal conditional distribution $Q_i^o(dx_i|x^{i-1},u^{i-1}) \in \mathcal{Q}(\mathcal{X}_i|\mathbb{K}_{0,i-1})$ having radius $R_i \in [0,2]$, for $i=0,1,\ldots,n$.

PROBLEM 3.6. *Given a nominal feedback controlled process of Definition 3.3 (a), an admissible policy set $\mathbf{\Pi}_{0,n-1}^{DF}$ and an ambiguity class $\mathbf{B}_{R_k}(Q_k^o)(x^{k-1},u^{k-1})$, k=0, ..., n of Definition 3.4 (a), find a $\pi^* \in \mathbf{\Pi}_{0,n-1}^{DF}$ and a sequence of stochastic kernels $Q_k^*(dx_k|x^{k-1},u^{k-1}) \in \mathbf{B}_{R_k}(Q_k^o)(x^{k-1},u^{k-1})$, $k=0,1,...,n$ which solve the following minimax optimization problem.*

$$\begin{aligned} J_{0,n}(\pi^*, Q_k^* : k=0,\ldots,n) = \inf_{\pi \in \mathbf{\Pi}_{0,n-1}^{DF}} \Big\{ \\ \sup_{\substack{Q_k(\cdot|x^{k-1},u^{k-1}) \in \mathbf{B}_{R_k}(Q_k^o)(x^{k-1},u^{k-1}) \\ k=0,1,\ldots,n}} \mathbb{E}_{\mathbf{Q}_\nu^\pi}\Big\{ \sum_{k=0}^{n-1} \alpha^k f_k(x_k^g,u_k^g) + \alpha^n h_n(x_n^g) \Big\} \Big\}. \end{aligned} \tag{3.6}$$

Next, we apply dynamic programming to characterize the solution of (3.6), by first addressing the maximization. Define the pay-off associated with the maximization problem

$$J_{0,n}(\pi, Q_k^* : k = 0, \ldots, n) \triangleq \sup_{\substack{Q_k(\cdot|x^{k-1},u^{k-1}) \in \mathbf{B}_{R_k}(Q_k^o)(x^{k-1},u^{k-1}) \\ k=0,1,\ldots,n}} J_{0,n}(\pi, Q_k : k = 0, \ldots, n).$$

For a given $\pi \in \mathbf{\Pi}_{0,n-1}^{DF}$, which defines $\{g_j : j = 0, \ldots, n-1\}$, and $\pi_{[k,m]} \equiv u_{[k,m]}^g$, denoting the restriction of policies in $[k, m]$, $0 \leq k \leq m \leq n-1$, define the conditional expectation taken over the events $\mathcal{G}_{0,j} \triangleq \sigma\{x_0^g, \ldots, x_j^g, u_0^g, \ldots, u_j^g\}$ maximized over the class $\mathbf{B}_{R_k}(Q_k^o)(x^{k-1}, u^{k-1})$, $k = j+1, \ldots, n$, as follows [7, 15]:

$$V_j(u_{[j,n-1]}^g, \mathcal{G}_{0,j}) \triangleq \sup_{\substack{Q_k(\cdot|x^{k-1},u^{k-1}) \in \mathbf{B}_{R_k}(Q_k^o)(x^{k-1},u^{k-1}) \\ k=j+1,\ldots,n}} \mathbb{E}_{\mathbf{Q}_\nu^\pi}\Big\{ \sum_{k=j}^{n-1} \alpha^k f_k(x_k^g, u_k^g) + \alpha^n h_n(x_n^g) | \mathcal{G}_{0,j} \Big\} \tag{3.7}$$

where $\mathbb{E}_{\mathbf{Q}_\nu^\pi}\{\cdot|\mathcal{G}_{0,j}\}$ denotes conditional expectation with respect to $\mathcal{G}_{0,j}$ calculated on the probability measure $\mathbf{Q}_\nu^\pi$. Then, $V_j(u_{[j,n-1]}^g, \mathcal{G}_{0,j})$ satisfies the following dynamic programming equation [15],

$$V_n(\mathcal{G}_{0,n}) = \alpha^n h_n(x_n^g) \tag{3.8}$$

$$V_j(u_{[j,n-1]}^g, \mathcal{G}_{0,j}) = \sup_{Q_{j+1}(\cdot|x^j,u^j) \in \mathbf{B}_{R_{j+1}}(Q_{j+1}^0)(x^j,u^j)} \Big\{ \mathbb{E}_{Q_{j+1}(\cdot|x^j,u^j)}\Big\{ \alpha^j f_j(x_j^g, u_j^g) + V_{j+1}(u_{[j+1,n-1]}^g, \mathcal{G}_{0,j+1}) \Big\} \Big\} \tag{3.9}$$

where $\mathbb{E}_{Q_{j+1}(\cdot|x^j,u^j)}\{\cdot\}$ denotes expectation with respect to $Q_{j+1}(dx_{j+1}|\mathbb{K}_{0,j})$.

Next, we present the dynamic programming recursion for the minimax problem. Let $V_j(\mathcal{G}_{0,j})$ represent the minimax pay-off on the future time horizon $\{j, j+1, ..., n\}$ at time $j \in \mathbb{N}_+^n$ defined by

$$V_j(\mathcal{G}_{0,j}) \triangleq \inf_{\pi \in \mathbf{\Pi}_{j,n-1}^{DF}} \sup_{\substack{Q_k(\cdot|x^{k-1},u^{k-1}) \in \mathbf{B}_{R_k}(Q_k^0)(x^{k-1},u^{k-1}) \\ k=j+1,\ldots,n}} \Big\{ \mathbb{E}_{\mathbf{Q}_\nu^\pi}\Big\{ \sum_{k=j}^{n-1} \alpha^k f_k(x_k^g, u_k^g) + \alpha^n h_n(x_n^g) | \mathcal{G}_{0,j} \Big\} \Big\} = \inf_{\pi \in \mathbf{\Pi}_{j,n-1}^{DF}} V_j(u_{[j,n-1]}^g, \mathcal{G}_{0,j}). \tag{3.10}$$

Then by reconditioning we obtain

$$V_j(\mathcal{G}_{0,j}) \triangleq \inf_{u \in \mathcal{U}_{ad}[j,n-1]} \sup_{\substack{Q_k(\cdot|x^{k-1},u^{k-1}) \in \mathbf{B}_{R_k}(Q_k^0)(x^{k-1},u^{k-1}) \\ k=j+1,\ldots,n}} \Big\{ \mathbb{E}_{\mathbf{Q}_\nu^\pi}\Big\{ \alpha^j f_j(x_j^g, u_j^g) + \mathbb{E}_{\mathbf{Q}_\nu^\pi}\Big\{ \sum_{k=j+1}^{n-1} \alpha^k f_k(x_k^g, u_k^g) + \alpha^n h_n(x_n^g) | \mathcal{G}_{0,j+1} \Big\} | \mathcal{G}_{0,j} \Big\} \Big\}. \tag{3.11}$$

Hence, we deduce the following dynamic programming recursion

$$V_n(\mathcal{G}_{0,n}) = \alpha^n h_n(x_n^g) \tag{3.12}$$

$$V_j(\mathcal{G}_{0,j}) \triangleq \inf_{u_j \in \mathcal{U}_j(x)} \sup_{Q_{j+1}(\cdot|x^j,u^j) \in \mathbf{B}_{R_{j+1}}(Q_{j+1}^0)(x^j,u^j)} \Big\{ \mathbb{E}_{Q_{j+1}(\cdot|x^j,u^j)}\Big\{ \alpha^j f_j(x_j^g, u_j^g) + V_{j+1}(\mathcal{G}_{0,j+1}) \Big\} \Big\}. \tag{3.13}$$

By applying the results of Section 2 to (3.12), (3.13) we obtain the following theorem.

THEOREM 3.7. *Suppose there exist an optimal policy for Problem 3.6, and assume* $V_{j+1}(\cdot):\mathcal{X}_{0,j+1}\times\mathcal{U}_{0,j}\longmapsto[0,\infty)$ *in (3.10) is bounded continuous in* $x\in\mathcal{X}_{j+1}$, $j=0,\ldots,n-1$.

1) The dynamic programming recursion is given by

$$V_n(\mathcal{G}_{0,n})=\alpha^n h_n(x_n^g) \tag{3.14}$$

$$\begin{aligned} V_j(\mathcal{G}_{0,j}) = \inf_{u_j\in\mathcal{U}_j(x)}\Big\{&\mathbb{E}_{Q^o_{j+1}}\Big(\alpha^j f_j(x_j^g,u_j^g)+V_{j+1}(\mathcal{G}_{0,j+1})|\mathcal{G}_{0,j}\Big)\\ &+\frac{R_j}{2}\Big(\sup_{x_{j+1}\in\mathcal{X}_{j+1}}V_{j+1}(\mathcal{G}_{0,j+1})-\inf_{x_{j+1}\in\mathcal{X}_{j+1}}V_{j+1}(\mathcal{G}_{0,j+1})\Big)\Big\}. \end{aligned} \tag{3.15}$$

Moreover,

$$V_j(\mathcal{G}_{0,j})=\inf_{u_j\in\mathcal{U}_j(x)}\mathbb{E}_{Q^*_{j+1}}\Big\{\alpha^j f_j(x_j^g,u_j^g)+V_{j+1}(\mathcal{G}_{0,j+1})|\mathcal{G}_{0,j}\Big\} \tag{3.16}$$

where, the optimal conditional distributions $\{Q_j^*: j=0,1,\ldots,n-1\}$ *are given by*

$$Q^*_{j+1}\left(\mathcal{X}^+_{j+1}|x^j,u^j\right)=Q^o_{j+1}(\mathcal{X}^+_{j+1}|x^j,u^j)+\frac{R_{j+1}}{2}\in[0,1],\ \ (x^j,u^j)\in\mathbb{K}_{0,j} \tag{3.17}$$

$$Q^*_{j+1}\left(\mathcal{X}^-_{j+1}|x^j,u^j\right)=Q^o_{j+1}(\mathcal{X}^-_{j+1}|x^j,u^j)-\frac{R_{j+1}}{2}\in[0,1],\ \ (x^j,u^j)\in\mathbb{K}_{0,j} \tag{3.18}$$

$$Q^*_{j+1}\left(A|x^j,u^j\right)=Q^o_{j+1}(A|x^j,u^j),\ \forall A\subseteq\mathcal{X}_{j+1}\backslash\mathcal{X}^+_{j+1}\cup\mathcal{X}^-_{j+1},\ \ (x^j,u^j)\in\mathbb{K}_{0,j} \tag{3.19}$$

and [2]

$$\mathcal{X}^+_{j+1}\triangleq\Big\{x_{j+1}\in\mathcal{X}_{j+1}:V_{j+1}(\mathcal{G}_{0,j+1})=\sup\Big\{V_{j+1}(\mathcal{G}_{0,j+1}):x_{j+1}\in\mathcal{X}_{j+1}\Big\}\Big\} \tag{3.20}$$

$$\mathcal{X}^-_{j+1}\triangleq\Big\{x_{j+1}\in\mathcal{X}_{j+1}:V_{j+1}(\mathcal{G}_{0,j+1})=\inf\Big\{V_{j+1}(\mathcal{G}_{0,j+1}):x_{j+1}\in\mathcal{X}_{j+1}\Big\}\Big\}. \tag{3.21}$$

2) The total pay-off is given by

$$J_{0,n}(\pi^*,Q_i^*:i=0,\ldots,n-1)=\sup_{Q_0(\cdot)\in\mathbf{B}_{R_0}(Q^o)}\mathbb{E}_{Q_0}\Big\{V_0(\mathcal{G}_{0,0})\Big\}. \tag{3.22}$$

Proof. 1) Consider (3.13) expressed in integral form

$$\begin{aligned} V_j(\mathcal{G}_{0,j}) = \inf_{u_j\in\mathcal{U}_j(x)}\Big\{&\alpha^j f_j(x_j,u_j)\\ &+\sup_{Q_{j+1}(\cdot|x^j,u^j)\in\mathbf{B}_{R_{j+1}}(Q^o_{j+1})(x^j,u^j)}\int V_{j+1}(\mathcal{G}_{0,j},z)Q_{j+1}(dz|x^j,u^j)\Big\}. \end{aligned} \tag{3.23}$$

By applying (2.5) we obtain (3.14), (3.15), while (3.17)-(3.21) follow as well.

2) By evaluating (3.10) at $j=0$ we obtain (3.22). This completes the derivation. □

By Theorem 3.7, the maximizing measure is given by (3.17)-(3.19), and it is a functional of the nominal measure. At this stage we cannot claim that the maximizing measure is Markovian, and hence the optimal strategy is not necessarily Markov. Therefore, the computation of optimal strategies using non-Markov nominal controlled processes is computationally intensive. Next, we restrict the minimax formulation to Markov controlled nominal processes.

[2] Note the notation Σ^0 and Σ_0 in Section 2 is identical to the notation $\mathcal{X}^+_{j+1}$ and $\mathcal{X}^-_{j+1}$, respectively.

3.3.2. Dynamic Programming for Finite Horizon D-MCM Subject to Ambiguity. Consider the Markov nominal controlled processes, based on Definition 3.3 (b), and define

$$V_j(u^g_{[j,n-1]}, \mathcal{G}_{0,j}) \triangleq \sup_{\substack{Q_k(\cdot|x^{k-1},u^{k-1})\in \mathbf{B}_{R_k}(Q^o_k)(x_{k-1},u_{k-1}) \\ k=j+1,\ldots,n}} \mathbb{E}_{\mathbf{Q}^\pi_\nu}\Big\{ \sum_{k=j}^{n-1} \alpha^k f_k(x^g_k, u^g_k) + \alpha^n h_n(x^g_n)|\mathcal{G}_{0,j}\Big\}.$$

In view of Section 2, specifically, the relation between the maximizing distribution and the nominal distribution (2.1)-(2.3), which also apply to conditional distributions, we deduce that the maximization conditional distribution $Q^*_i(dx_i|x^{i-1}, u^{i-1})$ is Markovian, hence $Q^*_i(dx_i|x^{i-1}, u^{i-1}) = Q^*_i(dx_i|x_{i-1}, u_{i-1})$, $\forall (x^{i-1}, u^{i-1}) \in \mathbb{K}_{0,i-1}$. This observation can be verified by checking expressions (3.17)-(3.19). Then we define

$$V_j(u^g, x) \triangleq \sup_{\substack{Q_k(\cdot|x_{k-1},u_{k-1})\in \mathbf{B}_{R_k}(Q^o_k)(x_{k-1},u_{k-1}) \\ k=j+1,\ldots,n}} \mathbb{E}_{\mathbf{Q}^\pi_\nu}\Big\{ \sum_{k=j}^{n-1} \alpha^k f_k(x^g_k, u^g_k) + \alpha^n h_n(x^g_n)|x\Big\}. \tag{3.24}$$

Utilizing the above observations we obtain the analog of Theorem 3.7 for finite horizon D-MCM, as follows.

Define the value function

$$V_j(x) = \inf_{\pi\in\mathbf{\Pi}^{DM}_{j,n-1}} \sup_{\substack{Q_k(\cdot|x_{k-1},u_{k-1})\in \mathbf{B}_{R_k}(Q^o_k)(x_{k-1},u_{k-1}) \\ k=j+1,\ldots,n}} \mathbb{E}_{\mathbf{Q}^\pi_\nu}\Big\{ \sum_{k=j}^{n-1} \alpha^k f_k(x^g_k, u^g_k) + \alpha^n h_n(x^g_n)|x\Big\}. \tag{3.25}$$

Then we obtain the following theorem.

THEOREM 3.8. *Suppose there exists an optimal policy for Problem 3.6, for the class of Markov Nominal Controlled Process of Definition 3.4 (b). Then the following hold:*

1) If the infimum over feedback strategies in (3.25) exists it is Markov $\pi \in \mathbf{\Pi}^{DM}_{0,n-1}$.

2) The value function $V_j(x)$ satisfies the dynamic programming recursion

$$V_n(x) = \alpha^n h_n(x), \quad x \in \mathcal{X}_n \tag{3.26}$$

$$V_j(x) = \inf_{u\in\mathcal{U}_j(x)} \sup_{Q_{j+1}(\cdot|x,u)\in\mathbf{B}_{R_{j+1}}(Q^o_{j+1})(x,u)} \mathbb{E}_{Q_{j+1}(\cdot|x,u)}\Big\{\alpha^j f_j(x,u) + V_{j+1}(x_{j+1})\Big\}, \quad x \in \mathcal{X}_j. \tag{3.27}$$

3) Assume $V_{j+1}(\cdot) : \mathcal{X}_{j+1} \to [0,\infty)$ is bounded continuous in $x \in \mathcal{X}_{j+1}$, $j = 0,\ldots,n-1$, then the dynamic programming recursion is given by

$$V_n(x) = \alpha^n h_n(x), \quad x \in \mathcal{X}_n \tag{3.28}$$

$$V_j(x) = \inf_{u\in\mathcal{U}_j(x)} \Big\{\alpha^j f_j(x,u) + \int_{\mathcal{X}_{j+1}} V_{j+1}(z) Q^o_{j+1}(dz|x,u) + \frac{R_j}{2}\Big(\sup_{z\in\mathcal{X}_{j+1}} V_{j+1}(z) - \inf_{z\in\mathcal{X}_{j+1}} V_{j+1}(z)\Big)\Big\}, \quad x \in \mathcal{X}_j. \tag{3.29}$$

Moreover,

$$V_j(x) = \inf_{u\in\mathcal{U}_j(x)} \mathbb{E}_{Q^*_{j+1}}\Big\{\alpha^j f_j(x^g_j, u^g_j) + V_{j+1}(x_{j+1})|x_j = x\Big\} \tag{3.30}$$

where the optimal conditional distribution $\{Q_j^*(\cdot|\cdot,\cdot): j=0,1,\ldots,n-1\}$ *is given by*

$$Q_{j+1}^*\left(\mathcal{X}_{j+1}^+|x_j,u_j\right)=Q_{j+1}^o(\mathcal{X}_{j+1}^+|x_j,u_j)+\frac{R_{j+1}}{2}\in[0,1],\ \ (x_j,u_j)\in\mathbb{K}_j \tag{3.31}$$

$$Q_{j+1}^*\left(\mathcal{X}_{j+1}^-|x_j,u_j\right)=Q_{j+1}^o(\mathcal{X}_{j+1}^-|x_j,u_j)-\frac{R_{j+1}}{2}\in[0,1],\ \ (x_j,u_j)\in\mathbb{K}_j \tag{3.32}$$

$$Q_{j+1}^*\left(A|x_j,u_j\right)=Q_{j+1}^o(A|x_j,u_j),\ \forall A\subseteq\mathcal{X}_{j+1}\backslash\mathcal{X}_{j+1}^+\cup\mathcal{X}_{j+1}^-,\ \ (x_j,u_j)\in\mathbb{K}_j \tag{3.33}$$

and

$$\mathcal{X}_{j+1}^+\triangleq\Big\{x_{j+1}\in\mathcal{X}_{j+1}:V_{j+1}(x_{j+1})=\sup\{V_{j+1}(x_{j+1}):x_{j+1}\in\mathcal{X}_{j+1}\}\Big\} \tag{3.34}$$

$$\mathcal{X}_{j+1}^-\triangleq\Big\{x_{j+1}\in\mathcal{X}_{j+1}:V_{j+1}(x_{j+1})=\inf\{V_{j+1}(x_{j+1}):x_{j+1}\in\mathcal{X}_{j+1}\}\Big\}. \tag{3.35}$$

4) The total minimax pay-off is

$$J_{0,n}(g^*,\{Q_i^*\}_{i=0}^n)=\sup_{Q_0(\cdot)\in\mathbf{B}_{R_0}(Q_0^o)}\mathbb{E}_{Q_0}\Big\{V_0(x_0)\Big\}. \tag{3.36}$$

Proof. 1) Since the nominal controlled process is Markov, from Theorem 3.7, (3.17)-(3.19) we deduce that the maximizing measure is also Markov. By the same arguments as in [15] we can show that if the infimum over $u\in\mathbf{\Pi}_{0,n-1}^{DF}$ in (3.25) exists, then it is Markov, and hence $u\in\mathbf{\Pi}_{0,n-1}^{DM}$.

2) By reconditioning we deduce that, the value function satisfies the dynamic programming equation (3.26), (3.27).

3) By definition, (3.27) is also equivalent to

$$V_j(x)=\inf_{u\in\mathcal{U}(x)}\Big\{\alpha^j f_j(x,u)+\sup_{Q_{j+1}(\cdot|x,u)\in\mathbf{B}_{R_{j+1}}(Q_{j+1}^o)(x,u)}\int_{\mathcal{X}_{j+1}}V_{j+1}(z)Q_{j+1}(dz|x,u)\Big\}.$$

Hence, by applying the results of Section 2 we obtain (3.28)-(3.33).

4) By evaluating (3.25) at $j=0$ we obtain (3.36). This completes the derivation. □

REMARK 3.9. *We make the following observations regarding Theorem 3.8.*
(a) The dynamic programming equation (3.28), (3.29) involves in its right hand side the oscillator seminorm of $V_{j+1}(\cdot)$.
(b) The dynamic programming recursion (3.28), (3.29) can be applied to controlled process with continuous alphabets and to controlled process with finite or countable alphabets, such as Markov Decision models.

Next, we show that for any $j\in\mathbb{N}^{n-1}$, the minimax pay-off $V_j(x)\equiv V_j^R(x)$ as a function of R_j is non-decreasing and concave.

LEMMA 3.10. *Suppose the conditions of Theorem 3.8 hold and in addition* $R_j=R$, $j=1,\ldots,n$. *The minimax pay-off* $V_j^R(x)\equiv V_j(x)$ *defined by (3.25) is a non-decreasing concave function of* R.

Proof. Consider two values for $R^1,R^2\in\mathbb{R}^+$ such that $0\leq R^1\leq R^2$. Since

$$\mathbf{B}_{R^1}(Q_k^o)(x_{k-1},u_{k-1})\subseteq\mathbf{B}_{R^2}(Q_k^o)(x_{k-1},u_{k-1})$$

then for every $Q_k(\cdot,x_{k-1},u_{k-1})\in\mathbf{B}_{R^1}(Q_k^o)(x_{k-1},u_{k-1})$ we have $Q_k(\cdot,x_{k-1},u_{k-1})\in B_{R^2}(Q_k^o)(x_{k-1},u_{k-1})$, $k=j+1,\ldots,n-1$. Hence, $V_j^{R^1}(x)\leq V_j^{R^2}(x)$ and thus, $V_j^R(x)$

is a non-decreasing function of $R \in \mathbb{R}^+$.

Next, for a fixed $\pi \in \Pi_{j,n-1}^{DM}$ consider two points (R^1, V_j^{π,R^1}), (R^2, V_j^{π,R^2}) such that $\{Q_k^1(\cdot|x_{k-1}, u_{k-1}) : k = j+1, \ldots, n\}$ achieves the supremum in (3.24) for R^1, and $\{Q_k^2(\cdot|x_{k-1}, u_{k-1}) : k = j+1, \ldots, n\}$ achieves the supremum in (3.24) for R^2. Then

$$||Q_k^1(\cdot|x_{k-1}, u_{k-1}) - Q_k^o(\cdot|x_{k-1}, u_{k-1})||_{TV} \leq R^1, \quad k = j+1, \ldots, n-1$$
$$||Q_k^2(\cdot|x_{k-1}, u_{k-1}) - Q_k^o(\cdot|x_{k-1}, u_{k-1})||_{TV} \leq R^2, \quad k = j+1, \ldots, n-1.$$

For any $\lambda \in (0,1)$ we have

$$\begin{aligned} &||\lambda Q_k^1(\cdot|x_{k-1}, u_{k-1}) + (1-\lambda)Q_k^2(\cdot|x_{k-1}, u_{k-1}) - Q_k^o(\cdot|x_{k-1}, u_{k-1})||_{TV} \\ &\leq \lambda||Q_k^1(\cdot|x_{k-1}, u_{k-1}) - Q_k^o(\cdot|x_{k-1}, u_{k-1})||_{TV} + (1-\lambda)||Q_k^2(\cdot|x_{k-1}, u_{k-1}) \\ &- Q_k^o(\cdot|x_{k-1}, u_{k-1})||_{TV} \leq \lambda R^1 + (1-\lambda)R^2, \qquad k = j+1, \ldots, n. \end{aligned} \tag{3.37}$$

Define $Q_k^*(\cdot|x_{k-1}, u_{k-1}) \triangleq \lambda Q_k^1(\cdot|x_{k-1}, u_{k-1}) + (1-\lambda)Q_k^2(\cdot|x_{k-1}, u_{k-1})$, $R = \lambda R^1 + (1-\lambda)R^2$. By (3.37), $Q_k^* \in \mathbf{B}_R(Q_k^o)(x_{k-1}, u_{k-1})$, $k = j+1, \ldots, n$. Define the unique probability measure

$$Q_{j+1,n}^*(dx^n||u^n) \triangleq \lambda \otimes_{k=j+1}^n Q_k^1(dx_k|x_{k-1}, u_{k-1}) + (1-\lambda) \otimes_{k=j+1}^n Q_k^2(dx_k|x_{k-1}, u_{k-1}).$$

Then,

$$V_j^{\pi,R}(x) \geq \int \Big(\sum_{k=j}^{n-1} f_k(x_k, u_k) + h_n(x_n) \Big) Q_{j+1,n}^*(dx^n||u^n)).$$

Hence,

$$\begin{aligned} V_j^{\pi,R}(x_j) &= \text{RHS of (3.24)} \\ &\geq \lambda \int \Big(\sum_{k=j}^{n-1} f_k(x_k, u_k) + h_n(x_n) \Big) \otimes_{k=j+1}^n Q_k^1(dx_k|x_{k-1}, u_{k-1}) \\ &\quad + (1-\lambda) \int \Big(\sum_{k=j}^{n-1} f_k(x_k, u_k) + h_n(x_n) \Big) \otimes_{k=j+1}^n Q_k^2(dx_k|x_{k-1}, u_{k-1}) \\ &= \lambda V_j^{\pi,R^1}(x_j) + (1-\lambda)V_j^{\pi,R^2}(x_j), \quad j = 0, \ldots, n-1. \end{aligned}$$

Hence, for any $\pi \in \mathbf{\Pi}_{j,n-1}^{DM}$, $V_j^{\pi,R}(x_j)$ is a concave function of R, and thus it is also concave for the $\pi \in \mathbf{\Pi}_{j,n-1}^{DM}$, which achieve the infimum in (3.25). □

This concavity property of the pay-off is also verified in the examples presented in Section 4.

REMARK 3.11. *The previous results apply to randomized strategies as well.*

3.4. Minimax Dynamic Programming for Infinite Horizon D-MCM Subject to Ambiguity. In this section, we consider the infinite horizon version of the finite horizon D-MCM, and we derive similar results. In addition, we show that the operator associated with the dynamic programming equation is contractive, and we introduce a new policy iteration algorithm.

Consider the problem of minimizing the finite horizon cost

$$\sup_{\substack{Q_k(\cdot|x,u) \in \mathbf{B}_{R_k}(Q_k^o)(x,u) \\ k=0,1,\ldots,n}} \mathbb{E}_{\mathbf{Q}_\nu^\pi} \Big\{ \sum_{j=0}^{n-1} \alpha^j f(x_j^g, u_j^g) \Big\} \tag{3.38}$$

with $0 < \alpha < 1$. By Theorem 3.8 the value function of (3.38), denoted by $V_j(x)$, $j = 0, \ldots, n$, $x \in \mathcal{X}_j$ satisfies the dynamic programming equations (3.28), (3.29) with $h_n = 0$, $R_j = R$, $\mathcal{X}_j = \mathcal{X}$, $\mathcal{U}_j = \mathcal{U}$, $\mathcal{U}_j(x) = \mathcal{U}(x)$ and $Q_j^o(\cdot|\cdot) = Q^o(\cdot|\cdot)$. Define $v_i(x) = \alpha^{i-n} V_{n-i}(x)$, where $0 \leq i \leq n$ is the time to go, (see [20]). Then,

$$v_0(x)=0 \tag{3.39}$$

$$v_i(x)=\inf_{u\in\mathcal{U}(x)}\Big\{f(x,u)+\alpha\int_{\mathcal{X}} v_{i-1}(z)Q^o(dz|x,u)+\alpha\frac{R}{2}\big(\sup_{z\in\mathcal{X}} v_{i-1}(z)-\inf_{z\in\mathcal{X}} v_{i-1}(z)\big)\Big\}. \tag{3.40}$$

In contrast with finite horizon case the one given by (3.39)-(3.40) proceeds from lower to higher values of indices i. The dynamic programming for the discounted cost

$$\mathbb{E}_{\mathbf{Q}_\nu^\pi}\Big\{\sum_{j=0}^{\infty}\alpha^j f(x_j^g,u_j^g)\Big\} \tag{3.41}$$

is given by

$$v_\infty(x) = \inf_{u\in\mathcal{U}(x)}\Big\{f(x,u)+\alpha\int_{\mathcal{X}} v_\infty(z)Q^o(dz|x,u)+\alpha\frac{R}{2}\Big(\sup_{z\in\mathcal{X}} v_\infty(z)-\inf_{z\in\mathcal{X}} v_\infty(z)\Big)\Big\}. \tag{3.42}$$

The maximizing conditional distribution is

$$Q^*\left(\mathcal{X}^+|x,u\right) = Q^o(\mathcal{X}^+|x,u) + \frac{R}{2} \in [0,1], \quad (x,u)\in\mathbb{K} \tag{3.43}$$

$$Q^*\left(\mathcal{X}^-|x,u\right) = Q^o(\mathcal{X}^-|x,u) - \frac{R}{2} \in [0,1], \quad (x,u)\in\mathbb{K} \tag{3.44}$$

$$Q^*\left(A|x,u\right) = Q^o(A|x,u), \quad \forall A\subseteq\mathcal{X}\setminus\mathcal{X}^+\cup\mathcal{X}^-, \quad (x,u)\in\mathbb{K} \tag{3.45}$$

where

$$\mathcal{X}^+ \triangleq \Big\{x\in\mathcal{X} : V(x) = \sup\{V(x) : x\in\mathcal{X}\}\Big\} \tag{3.46}$$

$$\mathcal{X}^- \triangleq \Big\{x\in\mathcal{X} : V(x) = \inf\{V(x) : x\in\mathcal{X}\}\Big\}. \tag{3.47}$$

Next, we show that the operator in the right hand side of (3.42) is contractive.

LEMMA 3.12. *Let L be the class of all measurable functions $V : \mathcal{X} \longrightarrow \mathbb{R}$, with finite norm $||V|| \triangleq \max_{x\in\mathcal{X}} |V(x)|$, and $T : L \longmapsto L$ defined by*

$$(TV)(x) = \inf_{u\in\mathcal{U}(x)}\Big\{f(x,u)+\alpha\int_{\mathcal{X}} V(z)Q^o(dz|x,u)+\alpha\frac{R}{2}\Big(\sup_{z\in\mathcal{X}} V(z)-\inf_{z\in\mathcal{X}} V(z)\Big)\Big\}. \tag{3.48}$$

If $V \in BC^+(\mathcal{X})$ and $\sup_{z\in\mathcal{X}} V(z)$, $\inf_{z\in\mathcal{X}} V(z)$ are finite, then T is a contraction.

Proof. For $V_1, V_2 \in L$,

$$\begin{aligned}
&(TV_1)(x) - (TV_2)(x) = \\
&\inf_{u\in\mathcal{U}(x)}\Big\{f(x,u)+\alpha\int_{\mathcal{X}} V_1(z)Q^o(dz|x,u)+\alpha\frac{R}{2}\Big(\sup_{z\in\mathcal{X}} V_1(z)-\inf_{z\in\mathcal{X}} V_1(z)\Big)\Big\} \\
&-\inf_{u\in\mathcal{U}(x)}\Big\{f(x,u)+\alpha\int_{\mathcal{X}} V_2(z)Q^o(dz|x,u)+\alpha\frac{R}{2}\Big(\sup_{z\in\mathcal{X}} V_2(z)-\inf_{z\in\mathcal{X}} V_2(z)\Big)\Big\}.
\end{aligned}$$

Let

$$v \triangleq \arg \inf_{u\in\mathcal{U}(x)} \Big\{ f(x,u) + \alpha \int_{\mathcal{X}} V_2(z) Q^o(dz|x,u) + \alpha \frac{R}{2} \Big(\sup_{z\in\mathcal{X}} V_2(z) - \inf_{z\in\mathcal{X}} V_2(z) \Big) \Big\}.$$

Then,

$$\begin{aligned}
&(TV_1)(x) - (TV_2)(x) \\
&= \inf_{u\in\mathcal{U}(x)} \Big\{ f(x,u) + \alpha \int_{\mathcal{X}} V_1(z) Q^o(dz|x,u) + \alpha \frac{R}{2} \Big(\sup_{z\in\mathcal{X}} V_1(z) - \inf_{z\in\mathcal{X}} V_1(z) \Big) \Big\} \\
&\qquad - \Big\{ f(x,v) + \alpha \int_{\mathcal{X}} V_2(z) Q^o(dz|x,v) + \alpha \frac{R}{2} \Big(\sup_{z\in\mathcal{X}} V_2(z) - \inf_{z\in\mathcal{X}} V_2(z) \Big) \Big\} \\
&\leq \Big\{ f(x,v) + \alpha \int_{\mathcal{X}} V_1(z) Q^o(dz|x,v) + \alpha \frac{R}{2} \Big(\sup_{z\in\mathcal{X}} V_1(z) - \inf_{z\in\mathcal{X}} V_1(z) \Big) \Big\} \\
&\qquad - \Big\{ f(x,v) + \alpha \int_{\mathcal{X}} V_2(z) Q^o(dz|x,v) + \alpha \frac{R}{2} \Big(\sup_{z\in\mathcal{X}} V_2(z) - \inf_{z\in\mathcal{X}} V_2(z) \Big) \Big\} \\
&\overset{(a)}{=} \Big\{ \alpha \int_{\mathcal{X}} V_1(z) Q^{V_1}(dz|x,v) \Big\} - \Big\{ \alpha \int_{\mathcal{X}} V_2(z) Q^o(dz|x,v) + \alpha \frac{R}{2} \Big(\sup_{z\in\mathcal{X}} V_2(z) - \inf_{z\in\mathcal{X}} V_2(z) \Big) \Big\} \\
&\overset{(b)}{\leq} \Big\{ \alpha \int_{\mathcal{X}} V_1(z) Q^{V_1}(dz|x,v) \Big\} - \Big\{ \alpha \int_{\mathcal{X}} V_2(z) Q^{V_1}(dz|x,v) \Big\} \\
&= \alpha \int_{\mathcal{X}} (V_1(z) - V_2(z))\, Q^{V_1}(dz|x|v) \leq \alpha \sup_{z\in\mathcal{X}} |V_1(z) - V_2(z)| = \alpha ||V_1 - V_2||
\end{aligned}$$

where (a) is obtained by applying (2.1), with $\ell \equiv \alpha V_1$, $\nu^*(\cdot) \equiv Q^{V_1}(\cdot|\cdot)$, $\mu(\cdot) \equiv Q^o(\cdot|\cdot)$, and (b) is obtained by first apply (2.1) as in (a) with Q^{V_2} and then replace Q^{V_2} by Q^{V_1} which is suboptimal hence, the upper bound. By reversing the roles of V_1 and V_2 we get $(TV_2)(x) - (TV_1)(x) \leq \alpha ||V_2 - V_1||$. Hence, $|(TV_1)(x) - (TV_2)(x)| \leq \alpha ||V_1 - V_2||$ for all $x \in \mathcal{X}$, and

$$||TV_1 - TV_2|| \triangleq \max_{x\in\mathcal{X}} |(TV_1)(x) - (TV_2)(x)| \leq \alpha ||V_1 - V_2||$$

which implies that the operator $T : L \longmapsto L$ is a contraction. □

Utilizing Lemma 3.12 we obtain the following theorem which is analogous to the classical result given in [20].

THEOREM 3.13. *Assume* $v_\infty \in BC^+(\mathcal{X})$ *and* $\sup_{z\in\mathcal{X}} v_\infty(z)$, $\inf_{z\in\mathcal{X}} v_\infty(z)$ *are finite.*
(1) The dynamic programming equation

$$v_\infty(x) = \inf_{u\in\mathcal{U}(x)} \Big\{ f(x,u) + \alpha \int_{\mathcal{X}} v_\infty(z) Q^o(dz|x,u) + \alpha \frac{R}{2} \Big(\sup_{z\in\mathcal{X}} v_\infty(z) - \inf_{z\in\mathcal{X}} v_\infty(z) \Big) \Big\}$$

has a unique solution.
(2) Moreover,

$$v_\infty(x) = \inf_{g\in\mathcal{U}(x)} \mathbb{E}_{Q^*} \Big\{ \sum_{j=0}^{\infty} \alpha^j f(x_j, u_j) | x_0 = x \Big\}.$$

(3) The mapping T *defined by*

$$(TV)(x) = \inf_{u\in\mathcal{U}(x)} \Big\{ f(x,u) + \alpha \int_{\mathcal{X}} V(z) Q^o(dz|x,u) + \alpha \frac{R}{2} \Big(\sup_{z\in\mathcal{X}} V(z) - \inf_{z\in\mathcal{X}} V(z) \Big) \Big\}$$

is a contraction mapping with respect to the norm $||V|| = \max_{x\in\mathcal{X}} |V(x)|$.

(4) For any V, $\lim_{n\to\infty} ||T^nV - v_\infty|| = 0$ *and so*

$$\lim_{n\longrightarrow\infty}(T^nV)(x) = v_\infty(x), \quad \textit{for all } \ x\in\mathcal{X}.$$

Proof. (1) Follows from [20] (Theorem 6.3.6, part (a)).

(2) We need to show that $v_\infty(x)$ is the minimum value of $\mathbb{E}_{Q^*}\Big\{\sum_{j=0}^{\infty}\alpha^j f(x_j,u_j)\Big\}$ starting in state $x_0 = x$. Recall that $0 \le f(x,u) \le M$ for all $x\in\mathcal{X}$, $u\in\mathcal{U}(x)$. Clearly, with $x_0 = x$ and for all n,

$$\inf_{g\in\mathcal{U}(x)}\mathbb{E}_{Q^*}\Big\{\sum_{j=0}^{\infty}\alpha^j f(x_j,u_j)\Big\} \ge \inf_{g\in\mathcal{U}(x)}\mathbb{E}_{Q^*}\Big\{\sum_{j=0}^{n-1}\alpha^j f(x_j,u_j)\Big\} = v_n(x).$$

Hence, $\inf_{g\in\mathcal{U}(x)}\mathbb{E}_{Q^*}\Big\{\sum_{j=0}^{\infty}\alpha^j f(x_j,u_j)\Big\} \ge \lim_{n\to\infty} v_n(x) = v_\infty(x)$. Conversely, for all n

$$\inf_{g\in\mathcal{U}(x)}\mathbb{E}_{Q^*}\Big\{\sum_{j=0}^{\infty}\alpha^j f(x_j,u_j)\Big\} \le \inf_{g\in\mathcal{U}(x)}\mathbb{E}_{Q^*}\Big\{\sum_{j=0}^{n-1}\alpha^j f(x_j,u_j)\Big\} + \sum_{j=n}^{\infty}\alpha^j M = v_n(x) + \frac{\alpha^n M}{1-\alpha}$$

and so

$$\inf_{g\in\mathcal{U}(x)}\mathbb{E}_{Q^*}\Big\{\sum_{j=0}^{\infty}\alpha^j f(x_j,u_j)\Big\} \le \lim_{n\to\infty}\Big[v_n(x) + \frac{\alpha^n M}{1-\alpha}\Big] = v_\infty(x).$$

Hence, $\inf_{g\in\mathcal{U}(x)}\mathbb{E}_{Q^*}\Big\{\sum_{j=0}^{\infty}\alpha^j f(x_j,u_j)\Big\} = v_\infty(x)$.

(3) This follows from Lemma 3.12.

(4) Follows from [20] (Theorem 6.3.6, part (b)). ∎

3.4.1. Policy Iteration Algorithm. Next, we present a modified version of the classical policy iteration algorithm [15]. From part 4 of Theorem 3.8, the policy improvement and policy evaluation steps of a policy iteration algorithm must be performed using the maximizing conditional distribution obtained under total variation distance ambiguity constraint. Hence, in addition to the classical case, in which the policy improvement and evaluation steps are performed using the nominal conditional distribution, here, under the assumption that $f(\cdot)$ is bounded and non-negative, by invoking the results developed in earlier sections we propose a modified algorithm which is expected to converge to a stationary policy in a finite number of iterations, since both state space $\mathcal{X}$ and control space $\mathcal{U}$ are finite sets, and that at each iteration a better stationary policy will be obtained.

First, we introduce some notation. Since the state space $\mathcal{X}$ is a finite set, with say, n elements, any function $V:\mathcal{X}\longrightarrow\mathbb{R}^n$ may be represented by vector in $\mathbb{R}^n$ defined by

$$V(x) \triangleq \begin{pmatrix} V(x_1) & \cdots & V(x_n)\end{pmatrix}^T \in \mathbb{R}^n.$$

Write $z \le y$, if $z(i) \le y(i)$, for $\forall i \in \mathbb{Z}^n \triangleq \{1,2,\ldots,n\}$; and $z < y$ if $z \le y$ and $z \ne y$. For a stationary control law g, let

$$f(g) = \begin{pmatrix} f(x_1,g(x_1)) & \cdots & f(x_n,g(x_n))\end{pmatrix}^T$$

and define each entry of the transition matrix $Q^o(g)\in\mathbb{R}^{n\times n}$ by $Q^o_{ij}(g) = Q^o(x_j|x_i, g(x_i)) \equiv Q^{g,o}(x_i|x_j)$. Rewrite (3.48) (with $\sup_{z\in\mathcal{X}} V(z)$ denoting componentwise supremum, and similarly for the infimum) as

$$TV = \min_{g\in\mathbb{R}^n}\Big\{f(g) + \alpha Q^o(g)V + \alpha\frac{R}{2}\Big\{\sup_{z\in\mathcal{X}} V(z) - \inf_{z\in\mathcal{X}} V(z)\Big\}\Big\}$$

which by Theorem 3.8 is equivalent to

$$TV = \min_{g\in\mathbb{R}^n}\Big\{f(g) + \alpha Q^*(g)V\Big\}$$

where $Q^*(g) \in \mathbb{R}^{n\times n}$ and is given by (3.43)-(3.45). Note that, the minimization is taken componentwise, i.e., $g(x_1)$ is the minimum of the first component of $f(g) + \alpha Q^*(g)V$ and so on. For each stationary policy g, define $T(g) : \mathbb{R}^n \longrightarrow \mathbb{R}^n$ by

$$T(g)V = f(g) + \alpha Q^*(g)V.$$

Then, $T(g)$ is a contraction mapping on the space of bounded continuous functions to itself, and from Theorem 3.13 it follows that

$$V(g) = T(g)V = f(g) + \alpha Q^*(g)V$$

has a unique solution $V(g) \in \mathbb{R}^n$. Next, we give the policy iteration algorithm.

ALGORITHM 3.14 (Policy Iteration). *Consider the notation above.*

Initialization. Let $m = 0$. *Solve the equation*

$$f(g_0) + \alpha Q^o(g_0)V_{Q^o}(g_0) = V_{Q^o}(g_0) \text{ for } V_{Q^o}(g_0) \in \mathbb{R}^n.$$

Identify the support sets using (3.46)-(3.47) and the analogue of Σ_k *of Section 2, and construct the matrix* $Q^*(g_0)$ *using (3.43)-(3.45). Solve the equation*

$$f(g_0) + \alpha Q^*(g_0)V_{Q^*}(g_0) = V_{Q^*}(g_0) \text{ for } V_{Q^*}(g_0) \in \mathbb{R}^n.$$

1. *For* $m = m + 1$ *while* $\min_{g\in\mathbb{R}^n}\Big\{f(g) + \alpha Q^*(g)V_{Q^*}(g_{m-1})\Big\} < V_{Q^*}(g_{m-1})$ *do:*

(a) (Policy Improvement) Let $g_m \in \mathbb{R}^n$ *be such that*

$$f(g_m) + \alpha Q^*(g_m)V_{Q^*}(g_{m-1}) = \min_{g\in\mathbb{R}^n}\Big\{f(g) + \alpha Q^*(g)V_{Q^*}(g_{m-1})\Big\}.$$

(b) (Policy Evaluation) Solve the following equation for $V_{Q^o}(g_m) \in \mathbb{R}^n$

$$f(g_m) + \alpha Q^o(g_m)V_{Q^o}(g_m) = V_{Q^o}(g_m).$$

Identify the support sets using (3.34)-(3.35), and construct the matrix $Q^*(g_m)$ *using (3.31)-(3.33). Solve the equation*

$$f(g_m) + \alpha Q^*(g_m)V_{Q^*}(g_m) = V_{Q^*}(g_m) \text{ for } V_{Q^*}(g_m) \in \mathbb{R}^n.$$

2. *Set* $g^* = g_m$.

In the next section, we illustrate through examples how the theoretical results obtained in preceding sections are applied.

4. Examples. In Section 4.1 we illustrate an application of the finite horizon minimax problem to the well-known machine replacement example, and in Section 4.2 we illustrate an application of the infinite horizon minimax problem for discounted cost by employing the policy iteration algorithm.

4.1. Finite Horizon MCM. Consider a machine replacement example inspired by [5]. Specifically, we have a machine that is either running or is broken down. If it runs throughout one week, it makes a profit of €100 for that week. If it fails during the week, the profit is zero for that week. If it is running at the start of the week and we perform preventive maintenance, the probability that it will fail during the week is 0.4. If we do not perform such maintenance, the probability of failure is 0.7. The maintenance cost is set at €20. When the machine is broken down at the start of the week, it may either be repaired at a cost of €40, in which case it will fail during the week with a probability of 0.4, or it may be replaced at a cost of €150 by a new machine that is guaranteed to run through its first week of operation. Assume that after $N>1$ weeks the machine, irrespective of its state, is scrapped with no cost.

The system dynamics is of the form $x_{k+1} = f_k(x_k, u_k, w_k)$, $k = 0, 1, \ldots, N-1$, where the state x_k is an element of a space $S_k = \{\mathbf{R}, \mathbf{B}\}$, $\mathbf{R}$ = machine running, $\mathbf{B}$ = machine broken, the control u_k is an element of a space $U_k(x_k)$, $U_k(\mathbf{R}) = \{m, nm\}$, m = maintenance, nm = no maintenance, $U_k(\mathbf{B}) = \{r, s\}$, r = repair, s = replace. The random disturbance has a nominal conditional distribution $w_k \sim \mu(\cdot|x_k, u_k)$.

Such a system can be described in terms of the discrete-time system equation $x_{k+1} = w_k$, where the nominal probability distribution of w_k is given by

$$\begin{aligned}
&\mu(w_k = \mathbf{R}|x_k = \mathbf{R}, u_k = m) = 0.6, && \mu(w_k = \mathbf{B}|x_k = \mathbf{R}, u_k = m) = 0.4,\\
&\mu(w_k = \mathbf{R}|x_k = \mathbf{R}, u_k = nm) = 0.3, && \mu(w_k = \mathbf{B}|x_k = \mathbf{R}, u_k = nm) = 0.7,\\
&\mu(w_k = \mathbf{R}|x_k = \mathbf{B}, u_k = r) = 0.6, && \mu(w_k = \mathbf{B}|x_k = \mathbf{B}, u_k = r) = 0.4,\\
&\mu(w_k = \mathbf{R}|x_k = \mathbf{B}, u_k = s) = 1, && \mu(w_k = \mathbf{B}|x_k = \mathbf{B}, u_k = s) = 0
\end{aligned}$$

and the input costs C_u are given by: if $u = m$ then $C_m =$ €20, if $u = nm$ then $C_{nm} =$ €0, if $u = r$ then $C_r =$ €40, and if $u = s$ then $C_s =$ €150. The cost per stage is $g_k(x_k, u_k, w_k) = C_{u_k}$ if $w_k = \mathbf{R}$, and $g_k(x_k, u_k, w_k) = C_{u_k} + 100$ if $w_k = \mathbf{B}$. Since it is assumed that after N weeks the machine, irrespective of its state, is scrapped without incurring any cost the terminal cost is $g_N(\mathbf{R}) = g_N(\mathbf{B}) = 0$.

The dynamic programming algorithm for the minimax problem subject to total variation distance uncertainty is given by

$$V_N(x_N) = 0 \tag{4.1}$$

$$\begin{aligned}
V_k(x_k) = &\min_{u_k \in U_k(x_k)} \max_{\nu(dw_k|x_k,u_k):||\nu(\cdot|x_k,u_k)-\mu(\cdot|x_k,u_k)||_{TV}\leq R} \Big\{ \\
&\mathbb{E}\Big\{ g_k(x_k, u_k, w_k) + V_{k+1}(f(x_k, u_k, w_k)) \Big\} \Big\} \qquad (4.2)\\
= &\min_{u_k \in U_k(x_k)} \max_{\nu(dw_k|x_k,u_k):||\nu(\cdot|x_k,u_k)-\mu(\cdot|x_k,u_k)||_{TV}\leq R} \mathbb{E}\Big\{ \ell_k(x_k, u_k, w_k) \Big\}
\end{aligned}$$

where $\ell_k(x_k, u_k, w_k) = g_k(x_k, u_k, w_k) + V_{k+1}(w_k)$, $k = 0, 1, \ldots, N-1$. To adress the maximization problem in (4.2), for each $k = 0, 1, \ldots, N-1$, $x_k \in \{\mathbf{R}, \mathbf{B}\}$ and $u_k \in \{m, nm, r, s\}$, define the maximum and minimum values of $\ell(x_k, u_k, w_k)$ by

$$\ell_{\max}(x_k, u_k) \triangleq \max_{w_k \in \{\mathbf{R},\mathbf{B}\}} \ell(x_k, u_k, w_k), \qquad \ell_{\min}(x_k, u_k) \triangleq \min_{w_k \in \{\mathbf{R},\mathbf{B}\}} \ell(x_k, u_k, w_k)$$

and its corresponding support sets by $\Sigma^0=\{w_k\in\{\mathbf{R},\mathbf{B}\}:\ell(x_k,u_k,w_k)=\ell_{\max}(x_k,u_k)\}$, and $\Sigma_0=\{w_k\in\{\mathbf{R},\mathbf{B}\}:\ell(x_k,u_k,w_k)=\ell_{\min}(x_k,u_k)\}$. By employing (2.5), the maximizing conditional probability distribution of the random parameter w_k is given by

$$\alpha = \min\Big(\frac{R}{2}, 1-\mu(\Sigma^0|x_k,u_k)\Big) \tag{4.3a}$$

$$\nu^*(\Sigma^0|x_k,u_k)=\mu(\Sigma^0|x_k,u_k)+\alpha, \qquad \nu^*(\Sigma_0|x_k,u_k)=\Big(\mu(\Sigma_0|x_k,u_k)-\alpha\Big)^+. \tag{4.3b}$$

Based on this formulation, the dynamic programming equation is given by

$$V_N(x_N)=0 \tag{4.4}$$

$$V_k(x_k)=\min_{u_k\in U_k(x_k)}\mathbb{E}_{\nu^*(\cdot|\cdot,\cdot)}\Big\{g_k(x_k,u_k,w_k)+V_{k+1}(f(x_k,u_k,w_k))\Big\}. \tag{4.5}$$

We assume that the planning horizon is $N=3$. The optimal cost-to-go and the optimal control policy, for each week and each possible state, as a function of $R\in[0,2]$ are illustrated in Figure 4.1. Clearly, Figure 4.1a depicts that the optimal cost-to-go is a non-decreasing concave function of R as stated in Lemma 3.10.

In addition, the optimum solution for two possible values of R and for each week results in optimal control policies as depicted in Table 4.1. By setting R=0, we choose to calculate the optimal control policy when the true conditional probability $\nu(\cdot|x_k,u_k)=\mu(\cdot|x_k,u_k),\ k$=0, 1, 2. This corresponds to the classical dynamic programming algorithm. By setting R=0.85, we choose to calculate the optimal control policy when the true conditional distribution $\nu(\cdot|x_k,u_k)\neq\mu(\cdot|x_k,u_k),\ k$=0, 1, 2. Taking into consideration the maximization (that is, by setting R>0) the dynamic programming algorithm results in optimal control policies which are more robust with respect to uncertainty, but with the sacrifice of low present and future costs. In cases in which we need to balance the desire for low costs with the undesirability of scenarios with high uncertainty, we must choose the appropriate value of R by using Figure 4.1a.

		Week.0		Week.1		Week.2	
	Stock	Cost-to-go	Optimal Policy	Cost-to-go	Optimal Policy	Cost-to-go	Optimal Policy
$R=0$	**R**	196	m	128	m	60	m
	B	216	r	148	r	80	r
$R=0.85$	**R**	340	m	221	m	100	nm
	B	360	r	241	r	122	r

Table 4.1: Dynamic Programming Algorithm Results

4.2. Infinite Horizon D-MCM. Here, we illustrate an application of the infinite horizon minimax problem for discounted cost, by considering the stochastic control system shown in Figure 4.2a, with state space $\mathcal{X}=\{1,2,3\}$ and control set $\mathcal{U}=\{u_1,u_2\}$.

Assume the nominal transition probabilities are given under controls u_1 and u_2 by

$$Q^o(u_1)=\frac{1}{9}\begin{pmatrix}3&1&5\\4&2&3\\1&6&2\end{pmatrix}, \qquad Q^o(u_2)=\frac{1}{9}\begin{pmatrix}1&2&6\\4&2&3\\4&1&4\end{pmatrix} \tag{4.6}$$

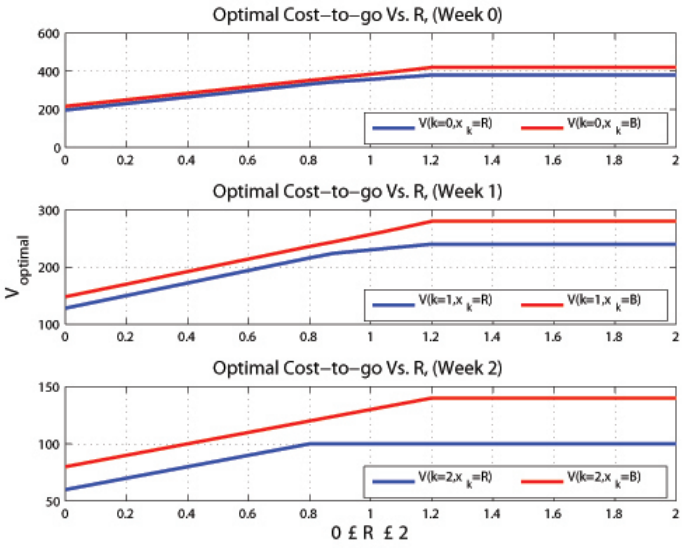

(a)

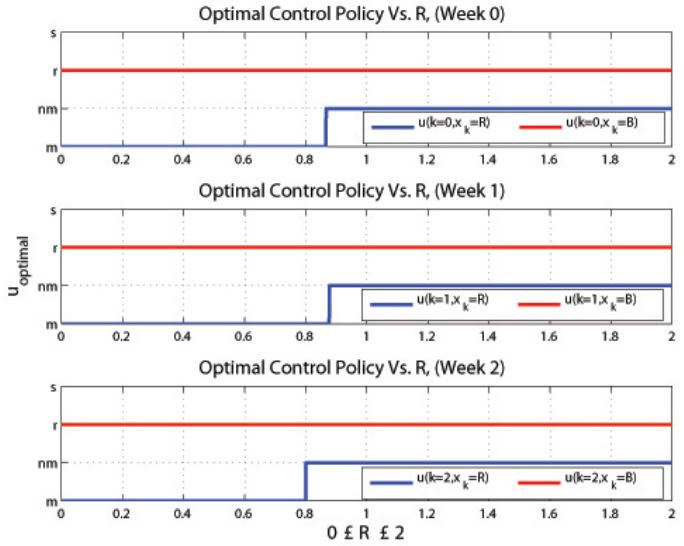

(b)

Fig. 4.1: (a) Optimal Cost-to-Go; (b) Optimal Control Policy ("m"= maintenance, "nm= no maintenace", "r=repair", "s=replace").

the discount factor is $\alpha = 0.9$, the total variation distance radius is $R = \frac{6}{9}$, and the cost function under each state and action is

$$f(1,u_1)=2,\ f(2,u_1)=1,\ f(3,u_1)=3,\ f(1,u_2)=0.5,\ f(2,u_2)=3,\ f(3,u_2)=0.$$

Using policy iteration of Section 3.3, with initial policies $g_0(1) = u_1$, $g_0(2) = u_2$, $g_0(3) = u_2$, the algorithm converge to the following optimal policy and value after two iterations.

$$g^* = g_2 \triangleq \begin{pmatrix} g_2(1) \\ g_2(2) \\ g_2(3) \end{pmatrix} = \begin{pmatrix} u_2 \\ u_1 \\ u_2 \end{pmatrix}, \qquad V_{Q^*}(g^*) = V_{Q^*}(g_2) \triangleq \begin{pmatrix} V_{Q^*}(1) \\ V_{Q^*}(2) \\ V_{Q^*}(3) \end{pmatrix} = \begin{pmatrix} 6.79 \\ 7.43 \\ 6.32 \end{pmatrix}.$$

Figure 4.2b depicts the optimal value functions for all possible values of R, and shows that, the value functions are non-decreasing and concave functions of R as stated in Lemma 3.10.

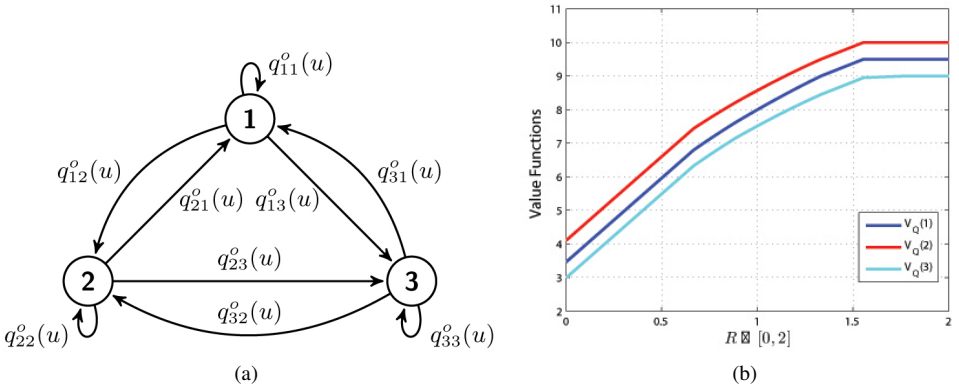

(a) (b)

Fig. 4.2: (a) Transition Probability Graph; (b) Optimal Value as a Function of Total Variation Parameter.

5. Conclusions. In this paper, we examined the optimality of stochastic control strategies via dynamic programming, when the ambiguity class is described by the total variation distance between the conditional distribution of the controlled process and the nominal conditional distribution. The problem is formulated using minimax strategies in which the control process seeks to minimize the pay-off while the controlled process seeks to maximize it over the total variation ambiguity class. By using concepts from signed measures a closed form expression of the maximizing measure is derived. It is then employed to obtain a new dynamic programming recursion which, in addition to the standard terms, includes the oscillator seminorm of the value function, while for the infinite horizon case a new discounted dynamic programming equation is obtained. It is shown that the dynamic programming operator is contractive, and a new policy iteration algorithm is developed for computing the optimal stochastic control strategies. Finally, we illustrate through examples the applications of our results.

REFERENCES

[1] N.U. AHMED, *Linear and nonlinear filtering for scientists and engineers*, World Scientific Publishers, Singapore, New Jersey, London, Hong Kong, 1999.

[2] J.S. BARAS AND M. RABI, *Maximum entropy models, dynamic games, and robust output feedback control for automata*, in Proceedings of the 44th IEEE Conference on Decision and Control, and the European Control Conference, Seville, Spain, Dec. 12–15, 2005.

[3] T.S. BASAR AND P. BERNHARD, *H-infinity optimal control and related minimax design problems: A dynamic game approach*, Collection Systèmes complexes, Birkhuser, 1995.
[4] A. BENSOUSSAN AND R. ELLIOT, *A finite dimensional risk-sensitive control problem*, SIAM J. Contol Optim., 33 (1995), pp. 1834–1846.
[5] D.P. BERTSEKAS, *Dynamic programming and optimal control*, Athena Scientific, 2005.
[6] D.P. BERTSEKAS AND S.E. SHREVE, *Stochastic optimal control: The discrete-time case*, Athena Scientific, 2007.
[7] P.E. CAINES, *Linear stochastic systems*, John Wiley & Sons, Inc., New York, 1988.
[8] C.D. CHARALAMBOUS AND J. HIBEY, *Minimum principle for partially observable nonlinear risk-sensitive control problems using measure-valued decompositions*, Stoch.Stoch.Rep., 57 (1996).
[9] C.D. CHARALAMBOUS AND F. REZAEI, *Stochastic uncertain systems subject to relative entropy constraints: Induced norms and monotonicity properties of minimax games*, IEEE Trans. Autom. Control, 52 (2007), pp. 647–663.
[10] C.D. CHARALAMBOUS, I. TZORTZIS, S. LOYKA, AND T. CHARALAMBOUS, *Extremum problems with total variation distance*, in 52nd IEEE Conference on Decision and Control, Florence, Italy, Dec. 10–13, 2013.
[11] R.J. ELLIOTT, L. AGGOUN, AND J.B. MOORE, *Hidden markov models: Estimation and control*, Springer, 1995.
[12] O. HERNANDEZ-LERMA AND J.B. LASSERRE, *Discrete-time markov control processes: Basic optimality criteria*, no. v. 1 in Applications of Mathematics Stochastic Modelling and Applied Probability, Springer Verlag, 1996.
[13] M. JAMES, J. BARAS, AND R. ELLIOT, *Risk-sensitive control and dynamic games for partially observed discrete-time nonlinear systems*, IEEE Trans. Autom. Control, 39 (1994), pp. 780–792.
[14] P.R. KUMAR AND J.H. VAN SCHUPPEN, *On the optimal control of stochastic systems with an exponential-of-integral performance index*, Journal of Mathematical Analysis and Applications, 80 (1981), pp. 312–332.
[15] P.R. KUMAR AND P. VARAIYA, *Stochastic systems: Estimation, identification, and adaptive control*, Prentice Hall, 1986.
[16] I.R. PETERSEN, M.R. JAMES, AND P. DUPUIS, *Minimax optimal control of stochastic uncertain systems with relative entropy constraints*, IEEE Trans. Autom. Control, 45 (2000), pp. 398–412.
[17] P.D. PRA, L. MENEGHINI, AND W.J. RUNGGALDIER, *Connections between stochastic control and dynamic games*, Math. Control Signals Systems, 9 (1996), pp. 303–326.
[18] F. REZAEI, C.D. CHARALAMBOUS, AND N.U. AHMED, *Optimal control of uncertain stochastic systems subject to total variation distance uncertainty*, SIAM Journal on Control and Optimization, 50 (2012), pp. 2683–2725.
[19] V.A. UGRINOVSKII AND I.R. PETERSEN, *Finite horizon minimax optimal control of stochastic partially observed time varying uncertain systems*, Math. Control Signals Systems, 12 (1999), pp. 1–23.
[20] J.H. VAN SCHUPPEN, *Mathematical control and system theory of discrete-time stochastic systems*, Preprint, 2010.
[21] P. WHITTLE, *A risk-sensitive maximum principle: The case of imperfect state observations*, IEEE Trans. Autom. Control, 36 (1991), pp. 793–801.